New developments in image-based characterization of coated particle nuclear fuel

Jeffery R. Price, Deniz Aykac, John D. Hunn, Andrew K. Kercher, Robert N. Morris

Oak Ridge National Laboratory, Oak Ridge, TN USA

ABSTRACT

We describe in this paper new developments in the characterization of coated particle nuclear fuel using optical microscopy and digital imaging. As in our previous work, we acquire optical imagery of the fuel pellets in two distinct manners that we refer to as shadow imaging and cross-sectional imaging. In shadow imaging, particles are collected in a single layer on an optically transparent dish and imaged using collimated back-lighting to measure outer surface characteristics only. In cross-sectional imaging, particles are mounted in acrylic epoxy and polished to near-center to reveal the inner coating layers for measurement. For shadow imaging, we describe a curvature-based metric that is computed from the particle boundary points in the FFT domain using a low-frequency parametric representation. We also describe how missing boundary points are approximated using band-limited interpolation so that the FFT can be applied. For cross-section imaging, we describe a new Bayesian-motivated segmentation scheme as well as a new technique to correct layer measurements for the fact that we cannot observe the true mid-plane of the approximately spherical particles.

Keywords: Image-based inspection, image segmentation, metrology, coated-particle nuclear fuel, TRISO fuel

1. INTRODUCTION

As noted in our previous paper,[1] the U.S. Department of Energy (DOE) Advanced Gas Reactor (AGR) Fuel Development and Qualification program was conceived to support near-term deployment of high-temperature, gas-cooled reactor technology and to establish a basis for the development of fuels suitable for very high-temperature, gas-cooled reactors. Continuing efforts in the first phase of the AGR program have concentrated on improving the production of high quality, silicon carbide-based, coated fuel particles (also known as ceramic-coated or TRISO fuel). Coated particle fuel development and characterization facilities have been established at Oak Ridge National Laboratory (ORNL) with the capability to characterize coated particles containing depleted, natural, and enriched uranium. This capability is required to support coating development work and predict the ultimate performance of the fuel particles. The goal of the work we present in this paper is to apply high-throughput, image-based measurements to enhance the efficiency and accuracy of the coating characterization.

An example image of a fully-coated particle is shown in Fig. 1, where the particle has been polished to near-center to reveal the structure of the coating layers. The image of Fig. 1 is an example of the cross-section images used in Section 3 below. From the center outward, the five layers of a fully-coated particle are the fuel *kernel*, a porous carbon *buffer*, inner pyrocarbon or *IPyC*, silicon carbide or *SiC*, and outer pyrocarbon or *OPyC*. Also evident in Fig. 1 is a shadow outside the OPyC; this shadow represents the excess particle since we cannot polish to the exact particle center (if particles are polished too near or past the center, the fuel kernels tend to dislodge).

We inspect particles using two types of imagery: shadow and cross-section. In shadow imaging, particles are collected in an optically transparent dish and imaged using collimated back-lighting. In these images, particles appear as silhouettes, as shown in Fig. 2, and the goal is to characterize the distribution of outer surface shapes based upon the perimeters of the observed silhouettes. In cross-section imaging, particles are mounted in a square array of acrylic epoxy, polished to near-center, and illuminated via both back-lighting and reflected light. This results in single-particle images, such as that shown in Fig. 1, and the goal is to measure the individual layer thicknesses at multiple points around the particle. Images are acquired using a Leica DMRX analytical, upright microscope with fluotar objectives. This microscope is fitted with a Leica DC500 camera, which employs

Correspondence: J. Price, pricejr@ornl.gov, 865-574-5743

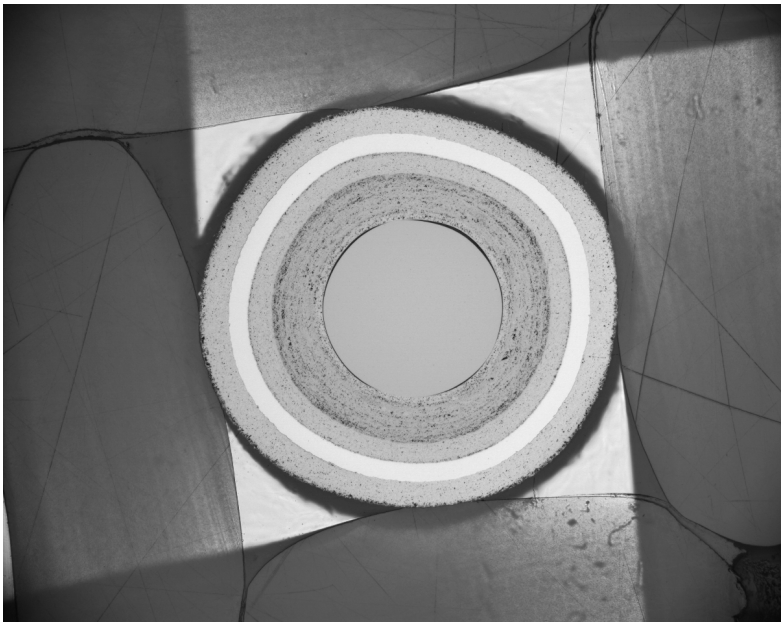

Figure 1. Cross-section view of a coated particle nuclear fuel pellet. Note the homogeneous kernel, which was a characteristic of the images in our earlier work. This 3900 × 3090 pixel image at 2.826 pixels/μm represents a region of approximately 1380μm × 1093μm.

a 1-megapixel Bayer mosaic chip and piezoelectric stepping to acquire 12-megapixel (3900 × 3090) images. For shadow imaging, image resolution is approximately 1.12 pixels/μm and for cross-section imaging, image resolution is approximately 2.83 pixels/μm. For fully-coated particles in which we are currently most interested, typical dimensions are approximately 175μm kernel radius, 100μm buffer thickness, 40μm IPyC thickness, 35μm SiC thickness, and 40μm OPyC thickness, resulting in particles of approximately 780μm in total diameter. Shadow imaging is typically performed on either fully coated particles or uncoated kernels.

The remainder of this paper is organized as follows. In Section 2, we discuss the analysis of shadow images, such as that shown in Fig. 2. We briefly review our earlier work and then describe more recent developments including band-limited interpolation to replace missing boundary points, a low-frequency sinusoidal approximation of the particle perimeter, and the computation and applicability of a curvature metric. In Section 3, we turn our attention to the analysis of cross-section imagery such as the example in Fig. 1. We describe a new, Bayesian-motivated approach for finding the layer boundaries. We also present an updated method to correct the observed layer thickness measurements for off-center measurement. We conclude in Section 4 with some brief closing remarks.

2. SHADOW IMAGE ANALYSIS

In this section, we describe measurements made using back-lit shadow imagery, such as the example shown in Fig. 2. Our goal is to characterize the distributions of various shape parameters (radius, eccentricity, curvature metric) of the approximately spherical particles by analyzing the silhouettes of many such particles that are presented in random orientations.

We previously described in detail[1] how individual particles are segmented and how their perimeter points are identified; we summarize that process here. Our first goal is simply to locate each particle in the image. We begin by subsampling the original image by a factor of four to speed processing. We next apply grayscale dilation to grow the bright background areas and effectively separate the particles from one another. We then threshold so that all dark regions (particles) are set to one and all background pixels are set to zero. The centroid, area, and eccentricity of each possible particle region are computed and used to discard any region that is too close to an image boundary or any likely non-particle regions (such as opaque regions of the particle holder as in the

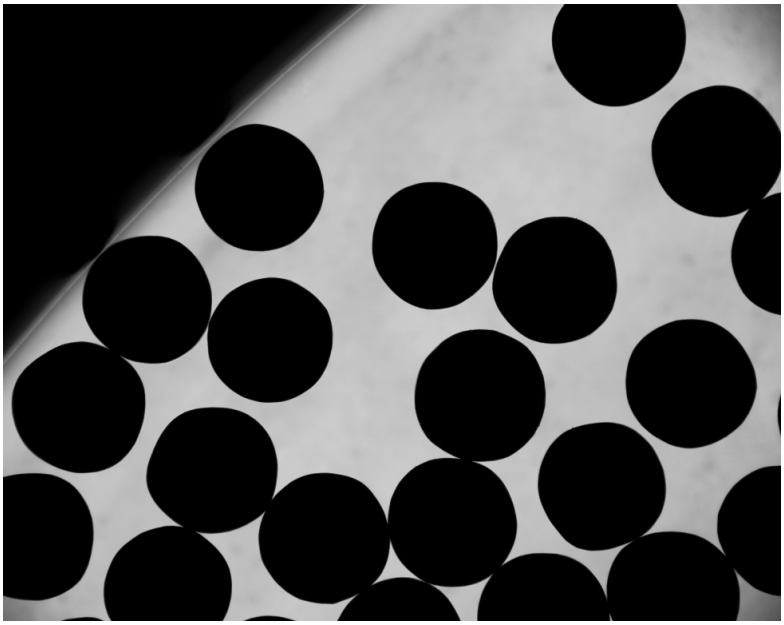

Figure 2. Example shadow image of multiple particles. This 3900 × 3090 pixel image at 1.1175 pixels/μm represents a region of approximately 3.49mm × 2.77mm.

upper left of Fig. 2). We next apply the distance transform and watershed segmentation to find the boundaries between different particle regions. An example result, based on the image shown earlier in Fig. 2, is shown in Fig. 3. We process each identified particle sequentially, using a full-resolution sub-image that is extracted from the original (i.e., not subsampled) image.

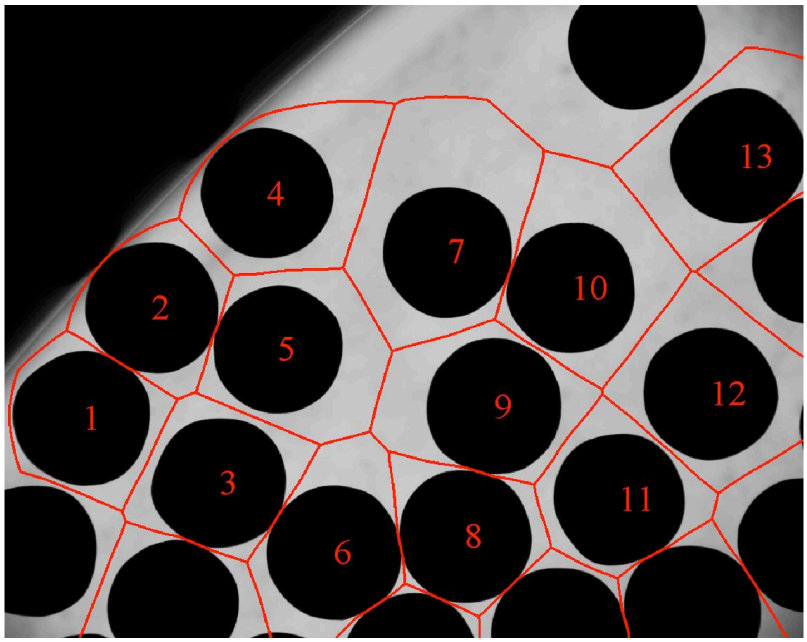

Figure 3. Segmentation of image from Fig. 2.

In each single-particle sub-image, we locate perimeter points using edge detection. We estimate the particle

center from these perimeter points using the Kasa circle fitting algorithm.[2,3] We next try to find N perimeter points spaced equally over $[0, 2\pi)$ about the estimated center, again using edge detection. Note that some perimeter points are not found due to abutting particles. To improve accuracy, we repeat the circle fit on these new perimeter points to update the center estimate. Finally, we seek N new perimeter points around the updated center. Our measurement process begins with these points.

2.1. Perimeter Representation

The perimeter points of a particle silhouette, found as described above, often capture small local variations due to surface roughness or small pieces of debris. These variations are not indicative of the particle shape we seek to characterize, can lead to higher measurement variance, and can significantly perturb sensitive curvature computations. We hence seek a smoother boundary. We represent the perimeter points as a parametric curve defined by $x(s)$ and $y(s)$ where the variable s represents the angular parameterization over $[0, 2\pi)$ about the estimated particle center. An ideally spherical particle would imply a circular silhouette with $x(s) = r\cos(s)$ and $y(s) = r\sin(s)$, where r is the radius. With this thought in mind, we compute FFTs of the sampled $x(s)$ and $y(s)$ and truncate higher frequency terms. To compute the FFTs, we need uniform sampling of $x(s)$ and $y(s)$ over the entire $[0, 2\pi)$ range. As mentioned previously, however, perimeter edge points are often missing due to abutting particles; see, for example, the particle labeled "8" in Fig. 3 and note how it abuts four other particles. This leads to missing samples on the perimeter curve, as can be seen in Fig. 4. To compute the FFTs then, we must first interpolate the missing samples of $x(s)$ and $y(s)$. To accomplish this in a manner consistent with our goal of low-frequency representation described above, we implement a bandlimited interpolation method suggested by Ferreira[4] and described as follows.

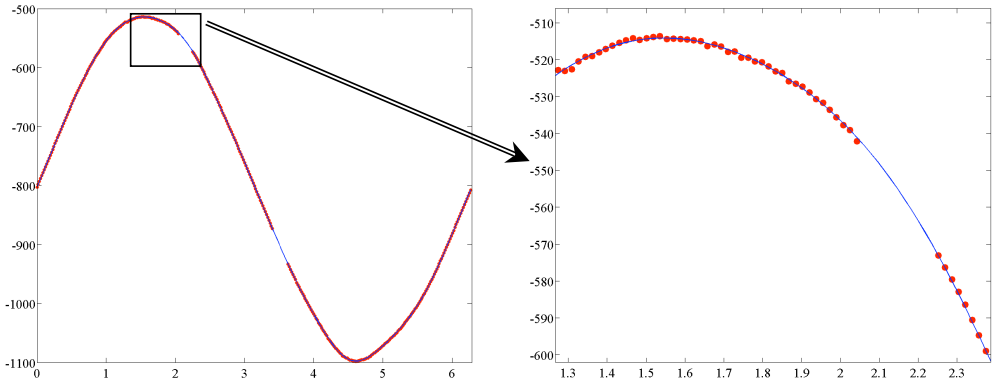

Figure 4. Plot of $x(s)$ for particle "3" from Fig. 3. The points indicate actual sample locations and the solid line represents the truncated FFT representation after bandlimited interpolation. Note the missing samples near 2.1 ($\approx 120°$) and 3.5 ($\approx 200°$) from abutting particles.

We define the vector $\mathbf{x} = [x_0\ x_1\ \ldots\ x_{N-1}]^T$, where $x_n = x(n\Delta s)$ and $\Delta s = 2\pi/N$. The N-point FFT of $\mathbf{x}$, which we define as $\mathbf{X}$, can be represented as a matrix multiply $\mathbf{X} = \mathrm{F}\mathbf{x}$, where the elements of the $N \times N$ matrix F are defined by $\mathrm{F}_{ik} = \exp(j2\pi ik/N)$. The inverse FFT can similarly be represented as $\mathbf{x} = \tilde{\mathrm{F}}\mathbf{X}$, where $\tilde{\mathrm{F}} = \frac{1}{N}\mathrm{F}^H$. If we assume that $\mathbf{x}$ is bandlimited, then some elements of $\mathbf{X}$ are zero; this can be expressed as $\mathbf{X} = \Gamma\mathbf{X}$, where the $N \times N$ matrix Γ is diagonal with only 1's or 0's on the diagonal. This leads to the following identity:

$$\mathbf{x} = \tilde{\mathrm{F}}\Gamma\mathrm{F}\mathbf{x} = \mathrm{B}\mathbf{x}. \tag{1}$$

Define $\mathcal{U}$ to be the set of indices for which $\mathbf{x}$ is unknown (i.e., the missing points) and $\mathcal{K}$ the set of indices for which $\mathbf{x}$ is known. Let the cardinality of $\mathcal{U}$ be P (hence the cardinality of $\mathcal{K}$ is $N - P$). For every $u \in \mathcal{U}$, we can write

$$x_u = \mathrm{B}_u\mathbf{x} = \sum_{i=0}^{N-1} \mathrm{B}_{ui}x_i = \sum_{i\in\mathcal{U}} \mathrm{B}_{ui}x_i + \sum_{i\in\mathcal{K}} \mathrm{B}_{ui}x_i, \tag{2}$$

where B_u represents row u of the matrix B. Concatenating the above expression for every $u \in \mathcal{U}$, we can write

$$\mathbf{x_u} = \mathrm{U}\mathbf{x_u} + \mathrm{K}\mathbf{x_k}, \tag{3}$$

where $\mathbf{x_u}$ is the $P \times 1$ vector corresponding to the unknown points of $\mathbf{x}$ that we seek, U is the $P \times P$ sub-matrix of B defined by taking the rows in $\mathcal{U}$ and columns in $\mathcal{U}$, K is the $P \times (N-P)$ matrix defined by taking the rows of B in $\mathcal{U}$ and columns in $\mathcal{K}$, and $\mathbf{x_k}$ is the $(N-P) \times 1$ vector corresponding to the known points (i.e., in $\mathcal{K}$), of $\mathbf{x}$. We can solve for the unknown points with

$$\mathbf{x_u} = (\mathrm{I} - \mathrm{U})^{-1}\mathrm{K}\mathbf{x_k}. \tag{4}$$

Recalling that we seek a low-frequency representation for $\mathbf{x}$, the matrix B is then computed as in Eq. (1) by setting all but the first $M+1$ and last M diagonal elements of Γ to one, where M defines the bandwidth we desire. Once we have estimated the missing samples of $x(s)$ and $y(s)$ in this manner, we simply take the N-point FFTs of $\mathbf{x}$ and $\mathbf{y}$, set to zero all but the first $M+1$ and last M coefficients, and then compute the inverse FFTs to yield the bandlimited representation of the silhouette boundary. An example is shown in Fig. 4.

2.2. Measurements

We compute several simple measurements for each particle from the bandlimited boundary points. These include mean radius, standard deviation of the radius, maximum radius, and minimum radius. Note that we only compute these measurements at points where we actually find the boundary (i.e., we do not use the interpolated missing points for measurement). We also compute the same measurements based upon diameters, where the diameter is defined by the distance between opposite boundary points. Historically, diameter aspect ratios (max diameter over min diameter) have been used as shape metrics for tolerance limit specifications. Aspect ratios, however, are gross descriptors of overall particle shape and are not necessarily correlated with potential structural failure except in extreme cases. Curvature-based metrics, on the other hand, can be used to identify sharp features that might serve as stress concentrators and lead more directly to possible failure.

The extrinsic curvature for a parametric curve $(x(s), y(s))$ can be expressed as

$$\kappa = \frac{x'y'' - y'x''}{\left((x')^2 + (y')^2\right)^{3/2}}, \tag{5}$$

where $x' \equiv dx/ds$. We analytically compute samples of the derivatives in Eq. (5) by applying frequency domain identities to the bandlimited FFT representations of $x(s)$ and $y(s)$. For example, let X_k for $k \in [0, N-1]$ represent the FFT coefficients of the N samples of $x(s)$ after bandlimited interpolation and zeroing of the higher-frequency terms as described in Section 2.1 above. We compute samples of $x'(s)$ by taking the inverse FFT of $X'_k = j\omega_k X_k$ where $\omega_k = [0, 1, \ldots, N/2, -N/2+1, -N/2+2, \ldots, -1]$ and samples of $x''(s)$ by taking the inverse FFT of $X''_k = (j\omega_k)^2 X_k$. With samples of the necessary derivatives computed in this way, we then compute samples of the curvature using Eq. (5). As a shape metric for a given particle, we compute the product of the maximum curvature and the radius at the point of maximum curvature; we refer to this metric as $R\kappa_{\max}$. It is a dimensionless quantity that is effectively scale invariant; it will equal 1.0 for a perfectly spherical particle (circular shadow), while a significantly larger number will indicate a sharp corner or protrusion. The $R\kappa_{\max}$ metric, under some significant simplifying assumptions, can be intuitively correlated with areas of high stress using membrane theory.[5]

We note that different shape metrics can emphasize differently the aspherical features present in particles. For example, diameter aspect ratio is an overall measure of particle eccentricity and can be influenced by up to four aspherical features corresponding to the opposing perimeter points of both the maximum and minimum diameter locations. Radius aspect ratio similarly measures eccentricity, but is affected by at most two aspherical features at the maximum radius (e.g., protrusion, corner) and minimum radius (e.g., facet, crater). When using diameter aspect ratio, the impact of a localized aspherical feature can sometimes be reduced because of the dependence upon the opposing perimeter point. Because of this, radius aspect ratio is typically a more sensitive measurement for quantifying localized aspherical deviations. The $R\kappa_{\max}$ metric is affected by only one aspherical feature: the sharpest corner. Many particles with high values of $R\kappa_{\max}$ are found acceptable according to both diameter and radius aspect ratios. As particle failure is often predicted by modeling, we are currently seeking to find the shape metrics that are most highly correlated with predicted particle failure.

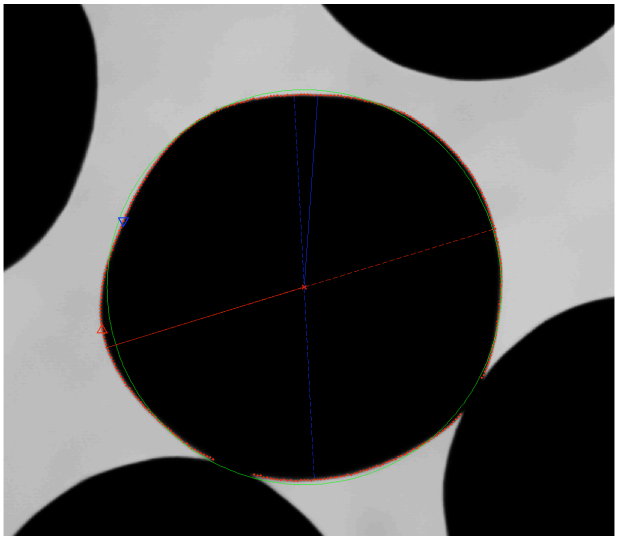

Figure 5. Example particle with a protrusion leading to a large value of $R\kappa_{\max} = 1.77$. The upward triangle indicates the point of highest curvature (the downward triangle indicates lowest curvature). The additional lines indicate maximal and minimal radii and diameters, where radii are solid lines and diameters are dashed lines.

3. CROSS-SECTION IMAGE ANALYSIS

Here we describe the analysis of cross-section images like the example shown earlier in Fig. 1. Our goal with this analysis is to characterize the thickness and uniformity of the interior layers to aid in evaluating and improving the coating processes. In preparation for imaging, the particles are mounted flat in a uniform mesh grid. The grid is then coated with transparent acrylic epoxy such that the surfaces of the particles are coplanar with the surface of the epoxy. This surface is polished away until cross-sections of the particles are visible, as shown earlier in the example of Fig. 1. The particles are not polished all the way to their center plane since the kernels might loosen, damage the buffer, and/or dislodge. Additionally, multiple particles, with some size variation, are mounted with the same bottom plane; hence, the center or mid-plane of one particle will not necessarily be the same as that for another particle of a different size. Note that we correct the layer thickness measurements for this off-center observation, as described in Section 3.2 below.

In our previous work,[1] we noted that the centers of the acquired images were always within the kernel region and that the kernel region was approximately homogeneous; this can be seen in the earlier example of Fig. 1. Necessary changes in the particle materials and image acquisition processes, however, have removed these helpful constraints from the image processing environment. An example of the images we are now analyzing is shown in Fig. 6. We previously used the homogeneity of the kernel, and the fact that the image center was inside the kernel, to perform a simple grayscale thresholding to find a rough estimate of the kernel region. We made an initial estimate of the particle center as the center of this approximate kernel region. Based upon this center estimate, we then sought the interface between the kernel and buffer, again making use of the kernel homogeneity compared to the roughness and generally lower gray values in the buffer region. As should be evident from Fig. 6, this approach will not work with the current images of interest. Instead, we have developed a simple program that sequentially presents each image of a selected directory and instructs a user to mouse-click one point near the kernel center and another point anywhere along the kernel-buffer interface. These two manually-entered points are then used as the initial estimate for the kernel center and approximate (*a priori*) location of the kernel-buffer interface.

Based on the manually-entered estimate of the kernel center, we compute an "unwrapped" image, like the example shown in Fig. 7, by resampling the original image on a polar grid point at one-degree angular resolution and one-pixel radial resolution. We search for the layer boundaries in this unwrapped image.

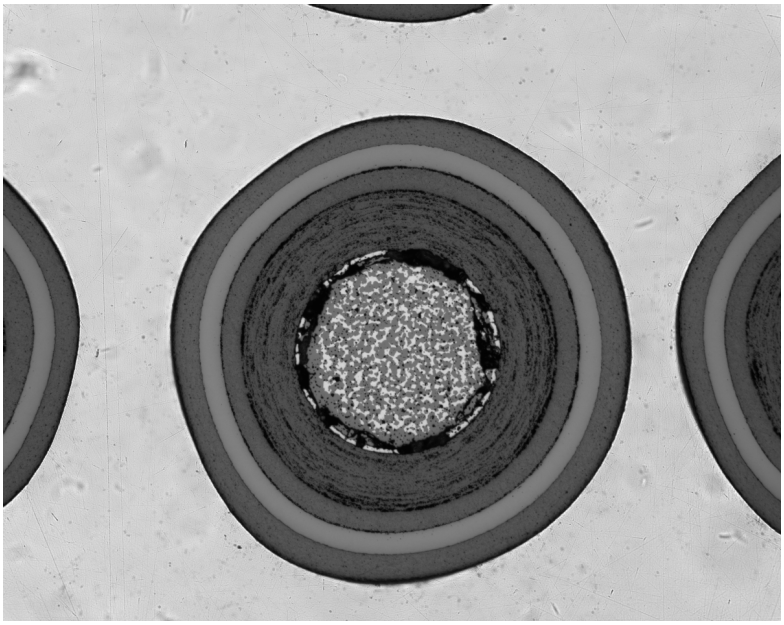

Figure 6. A more recent cross-section example image demonstrating the lack of homogeneity in the kernel region. We can also no longer guarantee that the image center lies within the kernel (although it does in this example).

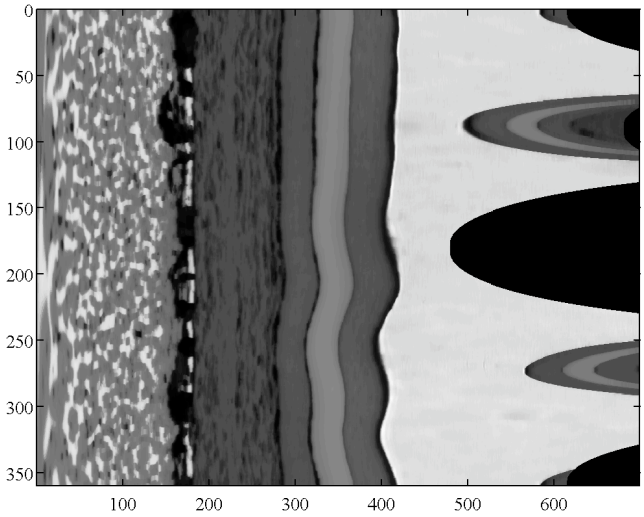

Figure 7. Unwrapped image computed by sampling the image of Fig. 6 on a polar grid about the particle center.

3.1. Finding the Boundaries

We previously employed[1] a rudimentary, deformable contour-like approach where a front of "marching points" moved from left-to-right across an unwrapped image. The movement of these points was controlled by sets of pragmatic, ad hoc rules based upon gray levels, gradients, and smoothness constraints. Here, we adopt a more formal, yet quite simple, approach using a Bayesian MAP-like methodology. We note that this new approach was conceived primarily as a means to initialize a subsequent (improved) deformable contour; this initialization, however, has proven accurate and robust enough to serve as the final segmentation.

Considering the boundary between two layers in an unwrapped cross-section image, we seek points on the

boundary, r_θ, such that the probability that r_θ is a boundary point is maximized. More formally, we seek

$$\max_{r_\theta}\Big(P(r_\theta|I) = \frac{P(I|r_\theta)P(r_\theta)}{P(I)}\Big), \tag{6}$$

where I represents the unwrapped image and θ represents the angle on the polar grid (vertical axis in Fig. 7). Since the image is fixed, we can ignore $P(I)$ in our computations. The $P(r_\theta)$ term represents *a priori* information on the boundary location. For a given layer, we model $P(r_\theta)$ as normally distributed about the design layer thickness or, in the case of the kernel-buffer interface, normally distributed about the distance between the manually-selected kernel center and interface point. The remaining term of interest from Eq. (6) is then $P(I|r_\theta)$, which is the probability of observing the image given that the point r_θ is a boundary point. We address this term as follows.

For each point in the unwrapped image, we consider the statistics of a region to its left, W_l, and a region to its right, W_r, as illustrated in Fig. 8. The triangular shapes of these regions are defined by boundary smoothness, which constrains the angles at the r_θ center point, and the expected layer dimensions, which constrain the horizontal extent of each region. For each pixel in an unwrapped image, we compute the mean and variance of the pixel values in both W_l and W_r. Note that these computations can be made using (four) simple filtering operations on the unwrapped image. We then compute the area under the overlap of the two Gaussians with these parameters. This area, P_{lr}, can be thought of as the probability of error in calling r_θ a boundary point, effectively quantifying how unlikely an image with boundary point r_θ is. Conversely, the quantity $1 - P_{lr}$ can be thought of as the probability of the image given that r_θ is a boundary point, i.e., $P(I|r_\theta)$.

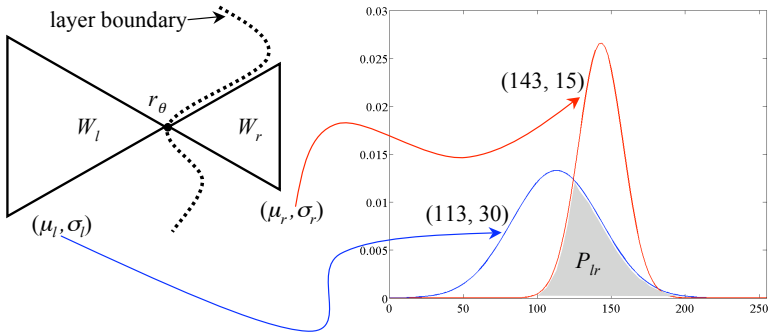

Figure 8. Region-based statistics used for boundary finding. For every pixel in the unwrapped image, r_θ, we consider regions to the left and right, W_l and W_r, respectively. The shapes of these regions are constrained by smoothness criteria and expected layer dimensions. We compute the means and standard deviations in these regions and then the amount of overlap, P_{lr}, between two normal distributions with these parameters. The quantity $1 - P_{lr}$ is used for $P(I|r_\theta)$. In this illustration, r_θ is shown on an example boundary (dotted line) that is at the limit of allowable smoothness.

We have now defined the two quantities of interest from Eq. (6) for each pixel in the unwrapped image. To find a complete layer boundary, we seek a set of points, $\{r_\theta\}$, that corresponds to a vertical path (i.e., 0 to 2π) in the unwrapped image such that the sum of all the $P(r_\theta|i)$ along this path is maximized relative to all of the possible paths. In other words, using the traditional log-likelihood expression, we seek to solve

$$\max_{\{r_\theta\}}\Big(\sum_\theta \log P(I|r_\theta) + \log P(r_\theta).\Big) \tag{7}$$

We accomplish this by first finding the maximum probability point in each row of the unwrapped image and begin a path with each of these points. Each of these paths is completed by progressing in a row-by-row (angular) manner. We examine a horizontal (radial) neighborhood about the previous point, where the horizontal extent of the neighborhood is constrained by our smoothness criteria, and add to the path the maximum probability point in the neighborhood. After all of the paths have been constructed in this manner, we select as the final path the one that satisfies Eq. (7). The final result of applying this segmentation approach for each layer of the unwrapped image in Fig. 7 is shown in Fig. 9.

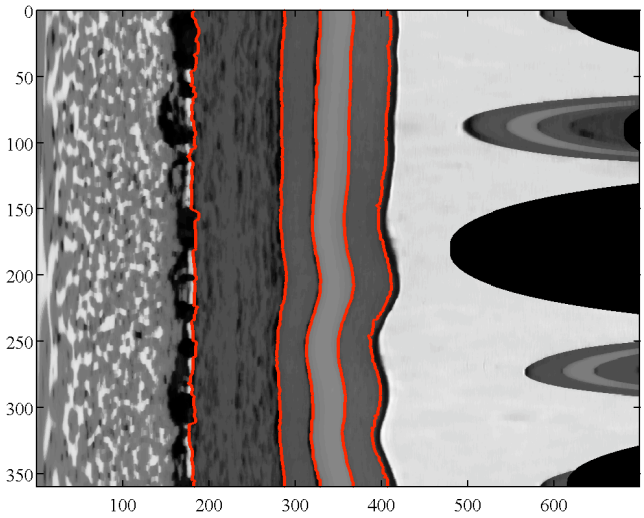

Figure 9. Result of segmentation applied to unwrapped image of Fig. 7.

3.2. Off-center Correction

As we mentioned in the beginning of Section 3, the cross-section plane observed in our images is generally not coincident with the true center plane of the particle. For this reason, the layer boundaries found as described cannot be used directly for accurate layer thickness measurements. This idea is represented in Fig. 10, which is a side view illustration of top-down cross-section imaging. As shown in the figure, radius of the layer boundary (dotted line) we measure from a cross-section image is represented as r_m, while the true center-plane radius is r. The polish-plane outer boundary of the outermost layer is represented by R_m while the true center-plane outer radius of the entire particle is R. The quantity p represents the amount of particle that is removed in the polishing process and is measured using a surface contact height gauge. With the measured quantities r_m, R_m, and p, we can show that the corrected layer position r is given by:

$$r = \left(r_m^2 + \left[\frac{1}{2p} (R_m^2 - p^2) \right]^2 \right)^{\frac{1}{2}} \tag{8}$$

We compute this correction for each boundary point of each layer. These corrected boundary points are used for all subsequent measurements (e.g., mean, min, and max layer thicknesses; standard deviation of layer thicknesses; etc.).

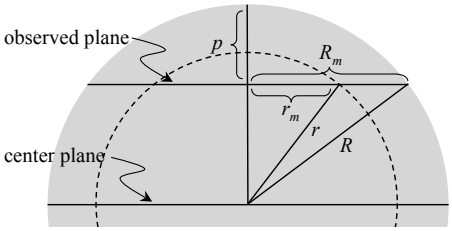

Figure 10. Side view illustration of cross-section imaging that shows the quantities used for correcting off-center observation. The measured quantity of interest is r_m while the true value is r.

4. CONCLUSIONS

We have described some recent developments in our use of image processing for the characterization of coated particle nuclear fuel. We discussed two types of coated particle image inspection, shadow imaging in Section 2 and cross-section imaging in Section 3. For shadow imaging, we described a new parametric boundary representation that made use of bandlimited interpolation and low-frequency sinusoidal representation. We also described the analytic computation of curvature in the FFT domain as well as a new curvature-based shape metric and discussed its applicability. For cross-section imaging, we described a MAP-like segmentation approach and discussed off-center correction for the observed layer dimensions. All software to date has been implemented and compiled using the MATLAB environment. Shadow image analysis is fully automated, while the cross-section analysis currently requires some minimal manual interaction (the selection of two landmark points in each image). The analysis programs have been used to analyze many thousands of particles to date and have proven very beneficial in improving our capabilities for manufacturing coated particle nuclear fuel.

ACKNOWLEDGMENTS

Prepared by Oak Ridge National Laboratory, managed by UT-Battelle, LLC, for the U.S. Department of Energy under Contract No. DE-AC05-00OR22725.

REFERENCES

1. J. R. Price and J. D. Hunn, "Optical inspection of coated particle nuclear fuel," in *Machine Vision Applications in Industrial Inspection XII*, **5303**, pp. 137–149, 2004.
2. I. Kasa, "A circle fitting procedure and its error analysis," *IEEE Transactions on Instrumentation and Measurement* **25**, pp. 8–14, March 1976.
3. C. Rusu, M. Tico, P. Kuosmanen, and E. Delp, "Classical geometrical approach to circle fitting – review and new developments," *Journal of Electronic Imaging* **12**, pp. 179–193, January 2003.
4. P. Ferreira, "Noniterative and fast iterative methods for interpolation and extrapolation," *IEEE Transactions on Signal Processing* **42**, pp. 3278–3282, November 1994.
5. P. Lowe, *Thin Shell Theory*, P. Noordhoff Ltd., 1964.