Shape from Interference Patterns

G. Cong and B. Parvin
Information and Computing Sciences Division
Lawrence Berkeley National Laboratory
Berkeley, CA 94720

Abstract

A unique imaging modality based on Equal Thickness Contours (ETC) has introduced a new opportunity for 3D shape reconstruction from multiple views. These ETCs are generated as a result of interference between transmitted and diffracted beams. We present a computational framework for representing each view of an object in terms of its *object thickness*, and then integrating these representations into a *3D surface* by algebraic reconstruction. In this framework, the object thickness is first derived from ideal contours and then extended to real data. For real data, the object thickness is inferred by grouping curve segments that correspond to points of second derivative maxima. At each step of the process, we use some form of regularization to ensure closeness to the original features, as well as neighborhood continuity. We apply our approach to images of a sub-micron crystal structure obtained through a holographic process.

1 Introduction

The problem of shape-from-X has been a central research topic in the computer vision community. These include–but are not limited to–shape from shading, texture, contour, color, etc. These techniques have been applied in a number of ways, from images obtained in controlled environments to natural outdoor scenes that may include more than one view. In this paper, we introduce an imaging modality, and the corresponding method for shape recovery, that has not yet been addressed by the computer vision community. This modality is based on equal thickness contour (ETC), obtained through a holographic process. One imaging source example is holographic electron microscopy of sub-micron crystal structures. Conventional electron microscopy presents projected images with little or no depth information. In contrast, electron holography–with coherent illumination–provides both magnitude and phase information that can be used to infer object thickness in terms of ETCs from each view of the sample. The holographic images contain interference fringes with spacings, in the best case down to less than an angstrom, where interference is between the transmitted and diffracted beams [9]. Similar imaging techniques include satellite radar interferometry, where the phase difference between two radar returns–at two different times from the same spatial location–is used for change detection [8]. It has been shown that minute geological changes (as a result of earthquake or movement in the earth's crust) can be recovered with this approach.

Interference patterns obtained from holographic image may have low contrast, be noisy, contain artifacts, and may have shading; as a result, it is difficult to compute closed contours from these fringe patterns. Figure 1 shows three views of a real crystal structure that will be used for shape recovery. A small angle of rotation between these views is indicated by the changes in the fringe patterns.

A critical step in recovering the shape from multiple views is the representation of each view of an object in terms of its thickness. This representation is derived from ETCs that are difficult to compute from real images. Hence, we first introduce a solution for recovering the thickness map for idealized perfect data and then show how this technique can be extended to real data. In the case of synthetic data, the problem reduces to "shape from sparse fringe contours," where a novel solution based on *equal importance criterion* is introduced. This constraint weighs each point in the image identically; it is expressed in a functional form and then minimized using variational principles. In the case of real images, we introduce a series of steps to recover the relevant features in the image, as shown in Figure 2. The contour representation (based on lines of second derivative maxima) provides the basis for obtaining object thickness at each point in the image. We have developed a regularized approach to computing object thickness, subject to continuity and smoothness constraints. Different views of an object are then merged through an extension of the algebraic reconstruction technique [6, 7].

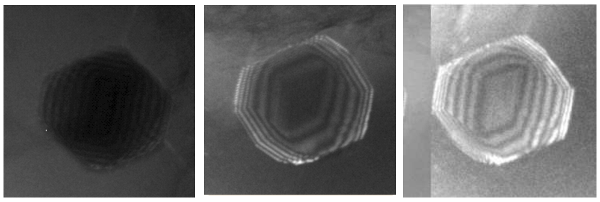

(a) (b) (c)

Figure 1: Three views of a cubeoctahedral object with a diameter of 100 nanometers.

*This work is supported by the Director, Office of Energy Research, Office of Computation and Technology Research, Mathematical, Information, and Computational Sciences Division under contract No. DE-AC03-76SF00098 with the University of California. The LBNL publication number is 41822. E-mail: gcong@george.lbl.gov or parvin@george.lbl.gov

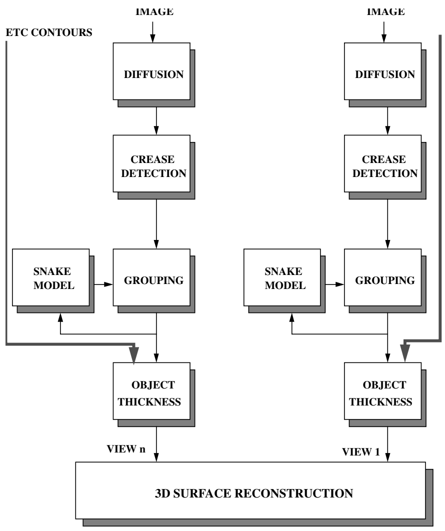

Figure 2: Protocol for recovering 3D shape from interference patterns

In the rest of this paper, we first introduce a recovery of thickness map from idealized contour data. In section 3, we will introduce steps required to recover the thickness map for real data. In section 4, the details of multiple view integration will be outlined, followed by concluding remarks in section 5.

2 Object thickness from ideal contours

The problem of shape from sparse data has been extensively studied in the computer vision and visualization community, as indicated in an excellent review by Bolle and Vemuri [4]. Although fringe contours are considered to be sparse, existing techniques are not applicable since the data is less dense, but well organized. Another related area to our research problem is surface reconstruction from slice contours that has been studied in computer graphics and medical image processing. However, these techniques assume that slices are very close to each other, which may not always be true. Furthermore, this technique [3] is based on Delaunay triangulation and thus scale dependent.

We propose to solve this ill-posed problem by enforcing the *Equal Importance Criterion*, which suggests that all points contribute equally to the shape reconstruction problem. Formally, the constraint states that surface height decreases linearly along the trajectory of its gradient, and as a result the problem becomes well posed. Since the solution for the system of partial differential equations (PDE) may not always exist, we redefine it as a variational problem to minimize the change in gradient over the region where the fringe contours are defined.

2.1 Fringe representation

Fringes may be defined mathematically in the following way. Let the 3D surface $f(x,y)$ be in a simply-connected region $(x,y) \in D$ with continuous first order derivative. Fringe contours $\mathcal{C}_i$ can be defined as:

$$\mathcal{C}_i = \{(x,y) | f(x,y) = i \cdot \triangle\xi\}, \qquad i = 1, ..., m \tag{1}$$

where $\triangle\xi$ is a constant and relates to the wave-length of the interferometry system. Without loss of generalities, we assume that $\triangle\xi = 1$. $f(x,y)$ is further constrained so that $\mathcal{C}_i$ has the following properties:

1. $\mathcal{C}_i$ is totally enclosed by $\mathcal{C}_j$ if and only if $i > j$.
2. The width of $\mathcal{C}_i$ is 0 (1 in the discrete case), *i.e.*, the surface does not have flat regions.

Let $R(\mathcal{C}_i, \mathcal{C}_j)$ be the doubly or multiply connected region enclosed by $\mathcal{C}_i$ and $\mathcal{C}_j$. The inverse problem can now be expressed as:

$$\begin{gathered} \textit{Given } \mathcal{C}_i, i = 1, ..., m, \textit{ find} \\ f(x,y) \textit{ where } (x,y) \in R(\mathcal{C}_1, \mathcal{C}_m), \textit{ such that } f(\mathcal{C}_i) = i. \end{gathered}$$

2.2 Equal Importance Criterion

First consider a simplified example in which we have only two contours $\mathcal{C}_1$ and $\mathcal{C}_2$. In $R(\mathcal{C}_1, \mathcal{C}_2)$, we want to construct a surface $f(x,y)$ such that $f(\mathcal{C}_1) = 1$, $f(\mathcal{C}_2) = 2$. Obviously in the absence of any constraints, infinitely many solutions exist. To constrain the problem, we assert that *every point between $\mathcal{C}_1$ and $\mathcal{C}_2$ is equally important and contributes similarly to the reconstruction process.* We call this the *Equal Importance Criterion.*

This constraint is formalized by requiring that the change in the gradient-magnitude along the gradient direction should be zero; that is:

$$\mathcal{J}_1(f) = \bigtriangledown(|\bigtriangledown f|) \cdot \frac{\bigtriangledown f}{|\bigtriangledown f|} = 0 \tag{2}$$

The above PDE implies that along each trajectory of the gradient of the surface, the magnitude of the gradient is a constant. In another words, the height decreases linearly from say 2 to 1. Thus, in the view of height, which is our only clue about the surface, all points are equally important to us. $\mathcal{J}_1$ can be reduced to:

$$\mathcal{J}_2(f) = f_x^2 f_{xx} + 2 f_x f_y f_{xy} + f_y^2 f_{yy} = 0 \tag{3}$$

where subscript indicates derivative, such as $f_x = \frac{\partial f}{\partial x}$, $f_{xy} = \frac{\partial^2 f}{\partial x \partial y}$. Thus, the problem becomes:

$$\begin{gathered} \textit{Find } \ f(x,y), \quad (x,y) \in R(\mathcal{C}_1, \mathcal{C}_2), \\ \textit{such that } \ \mathcal{J}_2(f) = 0, \quad f(\mathcal{C}_1) = 1, \quad f(\mathcal{C}_2) = 2. \end{gathered} \tag{4}$$

Equation (3) is called the "Infinity Laplacian", which has been studied in detail [1, 2]. Equation (4) has several important properties [1]:

1. There is at most one solution of the Dirichlet's problem for f. There is not always a smooth solution for any $\mathcal{C}_1$ and $\mathcal{C}_2$. However, if we redefine the "solution" in a suitable weak sense, then the solution does exist.
2. The trajectory of the vector field for $\bigtriangledown f$ is either a convex curve or a straight line.
3. A solution is infinitely differentiable if the trajectories of $\bigtriangledown f$ are convex curves.

All these properties are correct only in $R(C_1, C_2)$, where the value of f is known only on the boundaries.

2.3 Reconstruction of object thickness

In this section, a variational approach to shape reconstruction is proposed together with its numerical implementation.

It is natural to extend Equation (4) to $R(\mathcal{C}_1, \mathcal{C}_m)$:

$$\begin{gathered} \textit{Find } \ f(x,y), \quad (x,y) \in R(\mathcal{C}_1, \mathcal{C}_m), \\ \textit{such that } \ \mathcal{J}_2(f) = 0, \quad f(\mathcal{C}_i) = i, \quad i = 1, ..., m. \end{gathered} \tag{5}$$

However, this is an over-constrained problem, which may not always have an *exact* solution. We thus seek a solution for the following problem:

$$\begin{gathered} \textit{Minimize } \ \varepsilon^2 = \int\int_{R(\mathcal{C}_1, \mathcal{C}_m)} \mathcal{J}_1^2 \, dx dy, \\ \textit{Subject to}: \quad f(\mathcal{C}_i) = i, \quad i = 1, ...m. \end{gathered} \tag{6}$$

We can use an iterative approach: f can be solved iteratively by:

$$f^{v+1} = f^{v} - \sigma \delta f^{v} \tag{7}$$

where σ is the step size, v indicates the iteration number, and δf is the variation of f:

$$\delta f = -\frac{\partial \mathcal{F}_{f_x}}{\partial x} - \frac{\partial \mathcal{F}_{f_y}}{\partial y} + \frac{\partial^2 \mathcal{F}_{f_{xx}}}{\partial x^2} + \frac{\partial^2 \mathcal{F}_{f_{xy}}}{\partial x \partial y} + \frac{\partial^2 \mathcal{F}_{f_{yy}}}{\partial y^2} \tag{8}$$

2.4 Initialization

A good initialization is very important for deriving the solution to the desired local minima. This is accomplished by distance transform. Let $d_i(x,y)$ be the *Distance Transformation* of curve C_i, where it is negative inside C_i and positive outside of C_i. First, we represent $d_i(x,y)$ $i = 1,...,m$ as the Chamfer image of the fringes. Then $f(x,y)$, $(x,y) \in R(\mathcal{C}_i, \mathcal{C}_{i+1})$ is approximated as follows:

$$f(x,y) = \frac{d_i(x,y) \cdot (i+1) - d_{i+1}(x,y) \cdot i}{d_i(x,y) - d_{i+1}(x,y)} \tag{9}$$

2.5 Experimental results from ideal contours

Figure 3 shows the surface reconstructed from real data. The ETCs were interactively retrieved from the corresponding images. In all these images, $\sigma = 0.1$, and the algorithm converges in less than 100 iterations.

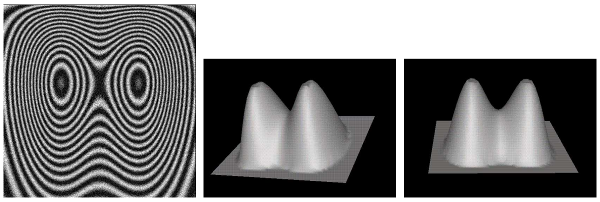

(a) (b) (c)

Figure 3: Reconstruction results: (a) Image used to generate perfect fringes interactively; (b)(c) two views of the corresponding reconstruction results.

3 Object thickness from real data

The original images are processed with a diffusion process to enhance the roof edges. Here, each point in the image is made smoother as a function of its inverse local Laplacian. Hence, when the Laplacian is large, no smoothing is performed. The maxima is then computed from the second directional derivative. These points are then linked and grouped using proximity, collinearity, and convexity. An example of grouping result is show in Figure 4. For more details, see our earlier paper [5].

3.1 Reconstruction of object thickness

The goal of reconstructing object thickness is to calculate a 2D image in which each point corresponds to the object thickness in a certain direction. The simple interpolation of section 2.4 is used to build the initial surfaces between the inner and outer fringes. Next, in the absence of closed contours, local contour segments provide the means for

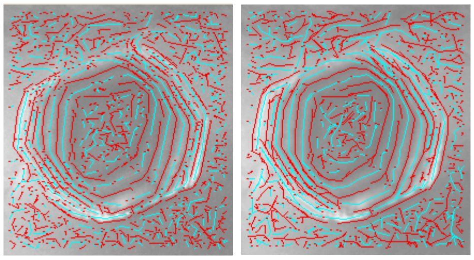

(a) (b)

Figure 4: Grouping results: (a) Computed contours from pointwise feature extraction; (b) result of grouping.

partial correction and updating in a local surface neighborhood. Adjacent contour segments are either parallel or intersect at a small angle. Consider trough T and ridge R contour segments that are adjacent to one another, and assume that from T to R the object thickness decreases by a constant, say 1, on the surface. If T and R are parallel with d distance apart (see Figure 5(a)) and the thickness decreases linearly, then the magnitude of gradient for points between T and R is given by: $|| \bigtriangledown f|| = \frac{1}{d}$. On the other hand, if the two segments are not parallel–see Figure 5(b)–then for each point p, there should be a curve passing through p that intersects T and R at p_1 and p_2, respectively. The tangent of this curve is in the direction of thickness gradient, which could also be obtained by tracking the gradient direction from p_1 to p_2. Since R and T are level set segments, it is easy to show that the normal to these segments at any point and the gradient of object thickness are in the same direction. When R and T are close to each other and the intersection angle between them is small, we can approximate a curve l, passing through p, by two line segments $pp_1' \perp T$, $pp_2' \perp R$, where $p_1'pp_2'$ is an approximation to p_1pp_2. Hence, $l \approx |p_1'p| + |p_2'p|$. These line segments are estimated from distance transforms (chamfer images), $d_t(x,y)$ and $d_r(x,y)$ that are computed from trough and ridge segments. Thus, $d_t = |p_1'p|$ and $d_r = |p_2'p|$, and the magnitude of the gradient is given by:

$$|| \bigtriangledown f(x,y)|| = \frac{1}{d_r(x,y) + d_t(x,y)} \tag{10}$$

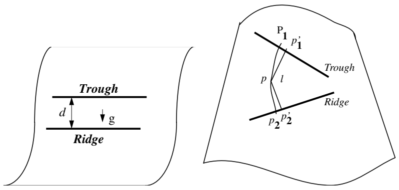

(a) (b)

Figure 5: Measure of gradient estimates from local contour segments: (a) T and R are parallel; (b) T and R are not parallel.

3.2 Regularized object thickness construction

The method for object thickness recovery (of section 2.3) is modified to accommodate the additional gradient thickness information of earlier section. The motivation is two fold:

1. $\mathcal{C}_i = i$ is no longer available; as a result, the actual thickness is unknown.
2. The computed thickness gradient is noisy and scale limited; as a result, some form of regularization is needed.

We have developed a regularized approach for constructing the object thickness $f(x,y)$. We define a suitable energy function to reflect the closeness to data and a smoothing term:

$$\epsilon^2 = \int\int (f_x - p)^2 + (f_y - q)^2 + \lambda \mathcal{J}_2^2 \, dx\, dy \qquad (11)$$

where f_x, f_y denote the x and y derivative of f; and p, q are the x and y direction gradient information computed earlier. Note that we do not apply additional constraints to show that the thickness along each segment is constant. However, the first term, $(f_x - p)^2 + (f_y - q)^2$ ensure that the gradient of each point on a segment is perpendicular to the segment, thus f is constant along the segment.

Following the Euler Equation, the surface f that minimizes (11) is:

$$(f_{xx} - p_x) + (f_{yy} - q_y) \; - \; \lambda \delta f = 0 \qquad (12)$$

where δf is the variation of $\mathcal{J}_2^2$ as defined in the previous section. This Equation is solved iteratively from an initial guess corresponding to the interpolation derived from the boundary conditions as indicated by Equation (9).

There is a certain ambiguity in filling up the region inside the inner contour as given by the initial condition. This ambiguity originates from the choice of planar fit or further interpolation from the inner contour to the center of the mass. This is a higher-level process and can only be resolved at the multiple-view integration. Figure 6 shows the reconstructed object thickness for each view of the object shown in Figure 1.

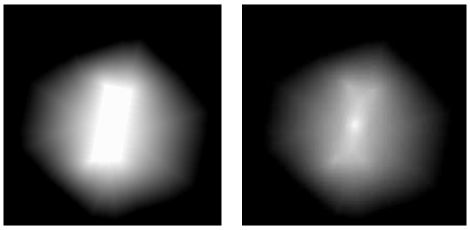

(a) (b)

Figure 6: Reconstructed thickness from the ETC in Figure 1. (a) planar fit; (b) interpolating inside inner contour.

4 3D shape recovery

In the preceding two sections, we developed an approach to recover the object thickness at each point along a certain direction. This is the thickness between front and back surfaces. In this section, we show how to integrate object thickness representation from different views into a 3D object. Our method extends the algebraic reconstruction techniques (ART) [6, 7]. However, we have a small number of views–three in our case–and the angles between these views are on the order of 7-8 degrees. The first step in 3D reconstruction is to align the object thickness representation bases in their location at center of mass. The 3D object boundary is then initialized as a polytope, which corresponds to the intersection of thickness extremes in 3D space. The 3D surface is then iteratively refined, based on object thickness information from each view. For simplicity, we present the details of our solution in 2D. Extension to 3D will be given in Appendix A. Let

- $g(i,j)$ be the 2D object image. $g(i,j), i = 0, ..., I, j = 0, ..., J$, is a discrete image that represents a binary object A and the background $\bar{A}$:

$$g(i,j) = \begin{cases} 1, & \text{if } (i,j) \in A \\ 0, & \text{if } (i,j) \in \bar{A} \end{cases} \qquad (13)$$

- $\mathcal{P}(\theta, k)$ be the computed thickness from angle θ at a point k as discussed in previous section,
- $p(\theta,k), -\frac{\pi}{4} \leq \theta \leq \frac{\pi}{4}$ be the thickness computed from $g(i,j)$. This thickness (as shown in Figure 7) is given by:

$$p(\theta, k) = \sum_{j=0}^{J} \frac{g(l_k, j)}{cos(\theta)} \qquad (14)$$

where (l_k, j) is the intersection point of line PN and the horizontal line $y = j$, and $g(l_k, j)$ is just a linear interpolation of the point around (l_k, j):

$$g(l_k, j) = g(i_1, j) * (i_2 - l_k) + g(i_2, j) * (l_k - i_1) \quad (15)$$

Thus, $p(\theta, k)$ is a linear combination of image pixels $g(i,j), i = 0, ..., I, j = 0, ..., J$.

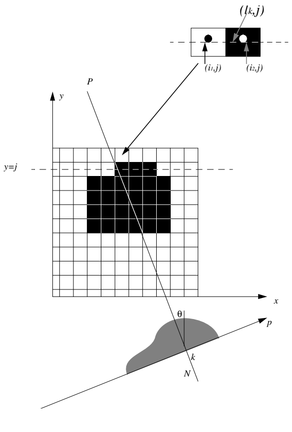

(a)

Figure 7: Projection of a 2D binary image on a line.

- $b(A), b(\bar{A})$ be the boundary points. The boundary point of the object is defined as:

$$b(A) := \{(i,j) | g(i,j) \cap [\bar{g}(i-1,j) \cup \bar{g}(i,j-1) \cup \bar{g}(i+1,j) \cup \bar{g}(i,j+1)] = 1\} \quad (16)$$

where $\cap$ and $\cup$ correspond to the logical *and* and *or* operations. $\bar{g}(i,j)$ is the inverse of $g(i,j)$.

Similarly, the boundary point of the background is defined as:

$$b(\bar{A}) := \{(i,j)|\bar{g}(i,j) \cap [g(i-1,j) \cup g(i,j-1) \cup g(i+1,j) \cup g(i,j+1)] = 1\} \tag{17}$$

- $\hat{p} \in b(\bar{A})$ of point $p \in b(A)$ be the neighborhood points, where $\hat{p}$ is defined over a 3-by-3 neighborhood of p in the background and is in the normal direction of the object's occluding contour at point p.

Assume that we have N thicknesses of a 2D object A along N directions $\mathcal{P}(\theta_n, k), n = 0, ..., N$, each thickness has M_n points. To reconstruct an $I \times J$ image $g(i,j)$ for the corresponding thicknesses, we define a cost function:

$$\epsilon^2 = \sum_{n=0}^{N} \sum_{m=0}^{M_n} (p(\theta_n, m) - \mathcal{P}(\theta_n, m))^2 + \lambda \sum_{(i,j) \in b(A)} (\mathcal{K}(i,j))^2 \tag{18}$$

where $\mathcal{K}$ is the product of the curvature of the boundary conditions and $| \bigtriangledown g|^3$. Since using the curvature would have complicated the optimization process, $\mathcal{K}$ is computed from the binary images as:

$$\mathcal{K} = g_x g_{yy} - 2 g_x g_y g_{xy} + g_y g_{xx} \tag{19}$$

The $\mathcal{K}$ term of the boundary point serves as a penalty term for smoothness, and the cost function is minimized with the gradient search scheme. There are two major differences with the ART method. First, we do not convert the 2D image into a 1D array vector, and second, we only calculate the derivative of ϵ with respect to points in $b(A)$ and $b(\bar{A})$. These points are updated according to the following criteria:

1. If $\frac{\partial \epsilon}{\partial g(p)} > 0$ and $\frac{\partial \epsilon}{\partial g(\bar{p})} > 0$, change $g(p)$ from 1 to 0.
2. If $\frac{\partial \epsilon}{\partial g(\bar{p})} < 0$ and $\frac{\partial \epsilon}{\partial g(\bar{p})} < 0$, change $g(\bar{p})$ from 0 to 1.

The results of shape recovery for the real object (of Figure 1) is shown in Figures 8.

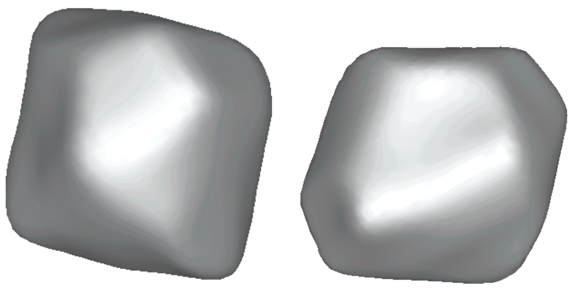

(a) (b)

Figure 8: 3D shape recovery from three views shown in Figure 1: (a) View1; (b) View2

5 Conclusion

In this paper, we introduced an imaging mode, in terms of equal thickness contours, that is being investigated for scientific studies. Next, we introduced a novel approach for shape recovery from multiple views and applied our approach to interference images of a sub-micron crystal structure. Our approach consists of three primary steps: (1) image representation in terms of contours corresponding to creases, (2) object thickness representation from contour data, and (3) shape recovery from object thickness obtained from each view.

A 3D algebraic reconstruction

The 3D extension of the algebraic reconstruction of Section 5 is accomplished by replacing $g(i,j)$ of Figure 7 with a cube $g(i,j,l)$, and the 1D projection $p(\theta, k)$ by a 2D projection $p(\vec{\theta}, k_1, k_2)$, where $\vec{\theta} = (\theta_1, \theta_2, \theta_3)$, and θ_1, θ_2, and θ_3 correspond to the angle between the projection direction with the x, y, and z axes. The projection can now be computed by:

$$p(\vec{\theta}, k_1, k_2) = \sum_{l=0}^{L} \frac{g(s_{k_1}, s_{k_2}, l)}{cos(\theta_3)} \tag{20}$$

where (s_{k_1}, s_{k_2}, l) is the point of intersection between the direction and plane $z = l$ and $-\frac{\pi}{4} \leq \theta_3 \leq \frac{\pi}{4}$. $g(s_{k_1}, s_{k_2}, l)$ is a bilinear interpolation in 2-by-2 neighborhood, and the projection is a linear combination of the image $g(i,j,l)$. The corresponding energy function is defined as:

$$\epsilon^2 = \sum_{n=0}^{N} \sum_{m_1=0}^{M_{1n}} \sum_{m_2=0}^{M_{2n}} (p(\vec{\theta}_n, m_1, m_2) - \mathcal{P}(\vec{\theta}_n, m_1, m_2))^2 + \lambda \sum_{(i,j,l) \in b(A)} (\mathcal{K}(i,j,l))^2 \tag{21}$$

where N is the number of projections of size M_{1n}-by-M_{2n} points, and $\mathcal{K}$ is the mean curvature of the boundary surface.

Acknowledgements: The authors thank (1) Drs. Ulrich Dahmen and Eric Johnson of the National Center for Electron Microscopy for providing the data; (2) Mr. Wes Bethel of the Visualization Laboratory for his help in using AVS; (3) Mr. Johnston –of the Information and Computing Sciences Division– for his encouragement in initiating this research.

References

[1] G. Aronsson. On the partial differential equation $u_x^2 u_{xx} + 2u_x u_y u_{xy} + u_y^2 u_{yy} = 0$. *Arikv för Matematik*, 7:133–151, 1968.

[2] G. Aronsson. On certain singular solutions of the partial differential equation $u_x^2 u_{xx} + 2u_x u_y u_{xy} + u_y^2 u_{yy} = 0$. *Manuscripta Math*, 41:133–151, 1981.

[3] J. Boissonnat. Shape reconstruction from planar cross sections. *Computer Vision, Graphics, and Image Processing*, 44(1):1–29, 1988.

[4] R.M. Bolle and B.C. Vemuri. On three-dimensional surface reconstruction methods. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 13(1):1–13, January 1991.

[5] G. Cong and B. Parvin. Shape from equal thickness contours. *Proceedings of the Conference on Computer Vision and Pattern Recognition*, June 1998.

[6] R. Gordon and G. Herman. Reconstruction of picture from their projections. *Communications of the ACM*, 14:759–768, 1971.

[7] R. Kashyap and M. Mittal. Picture reconstruction for projections. *IEEE Transactions on Computers*, 24:915–923, 1975.

[8] D. Massonnet. Satellite radar interferometry. *Scientific American*, 276(2):46–53, February 1997.

[9] A. Tonomura. *Progress in holographic interference electron microscopy*. North Holland, 1995.