3-D reconstruction of opaque objects from Fourier intensity data

M.F. Reiley[*†], *R.G. Paxman*[*†], *J.R. Fienup*[*†], *K.W. Gleichman*[†], and *J.C. Marron*[†]

[*] Advanced Modular Power Systems
4667 Freedom Drive
Ann Arbor, MI 48108-9784

[†] ERIM International
P.O. Box 134008
Ann Arbor, Michigan 48113-4008

Abstract

Three-dimensional imaging provides important profile information not available with conventional two-dimensional image products. Profile information can be extremely valuable for industrial-inspection and remote target-characterization applications. In this paper, we discuss a novel imaging modality, called PROCLAIM, that utilizes the powerful constraint that opaque objects can be described by a two-dimensional surface embedded in three-dimensional space. Far-field Fourier intensity measurements are collected by flood-illuminating an object with a frequency-tunable laser and direct detecting the backscattered signal with a lensless sensor. This technique allows for precise, non-contact surface measurements, without the stringent coherence and mechanical stability requirements of related interferometric techniques. We present reconstruction results from simulated data and from laboratory measurements.

Keywords: 3-D imaging, laser imaging, phase retrieval, profile retrieval, opacity, support bounds

1. Introduction

Three-dimensional (3-D) imaging provides profile (surface shape) information not available with conventional 2-D imaging. Profile information is known to be very useful for problems such as discrimination, identification, and surface metrology. Two broad classes of problems for which 3-D imaging can be used to significant advantage are: (1) industrial inspection and (2) remote target characterization.

Modern manufacturing processes are capable of producing parts to a high degree of mechanical precision. Industrial inspection techniques must determine whether these parts are in compliance with the strict tolerances imposed by design specifications. Contact surface metrology is often accomplished by Coordinate Measuring Machines (CMMs), providing accurate point-by-point measurements, which are valuable for spot-checks but are far too time consuming for complete or routine surface characterization. Clearly, such an approach is not well suited to the production environment. Optical surface metrology is a candidate technology that offers the potential for high-speed, low-cost measurements of surface relief.

The general class of problems that we refer to as remote target characterization includes detection, discrimination, and identification of remote targets such as missiles, aircraft, satellites, and ground targets. Such tasks are often performed using passive 2-D (cross-range) imaging sensors in infrared and visible spectral regions. Unfortunately, aberrations induced by the intervening atmosphere can severely limit the resolution achieved. Even when turbulence problems are mitigated [1], the image products are 2-D. The added value of profile information for performing discrimination and identification of targets is substantial [2]; hence the motivation for 3-D imaging of these targets.

In this paper we describe a novel 3-D imaging modality called *Phase Retrieval with an Opacity Constraint in LAser IMaging* (PROCLAIM). PROCLAIM is an active-illumination imaging method that can provide low-cost, fine-resolution profile information. PROCLAIM uses direct detection of the far-field Fourier intensity without intervening optics. As a consequence, the required detection hardware is exceptionally simple, and mechanical stability and coherence requirements are substantially relaxed relative to methods that require the detection of the complex optical field. In addition, PROCLAIM measurements are not impacted by turbulence-induced phase errors in upward-looking scenarios. The price that must be paid for these benefits is increased computer processing.

In the following section we review the PROCLAIM imaging sensor concept. We then present an end-to-end image-reconstruction example using simulated PROCLAIM data. The collection of real laboratory data is subsequently described. Finally, we present results from processed laboratory data.

2. Review of PROCLAIM

Image-reconstruction algorithms often use *a priori* object constraints to great advantage. For example, object support and nonnegativity have been extensively explored for use with ill-posed or ill-conditioned inverse problems. Another powerful constraint that has only recently received attention is object opacity [3,4,5]. An opaque object is one that exhibits only surface scattering and no volume scattering. The reflectivity function for an opaque object is confined to a 2-D (possibly curved or even discontinuous) surface, embedded in a 3-D space. The opacity constraint is a "quality of support" constraint – the actual location of the support is not given, although the object is known to be confined to a 2-D curved surface. This constraint promises to be very powerful since it greatly reduces the class of feasible objects from which to choose an estimate. Moreover, there are many imaging applications in which the objects will be known with confidence to be opaque. The constraint is invalid for objects with distributed volume scatterers such as translucent or fog-like objects. The use of opacity is most compelling in the context of 3-D imaging. Indeed, opacity is a fundamental component of the PROCLAIM imaging-sensor concept.

PROCLAIM is an active-illumination imaging method that requires the flood illumination of an opaque object with a frequency-tunable laser. The reflected radiation at a single frequency will create a speckle pattern in the far-field. The intensity of this far-field speckle pattern is directly detected with an array of detectors and without intervening optics. Typically, the illuminating laser will step through several frequencies so that a separate cross-range speckle intensity pattern is collected for each of multiple frequencies. Properly formatted, these data correspond to the modulus squared of the Fourier transform of the object's 3-D complex reflectivity function [6]. Note that a Fourier transform of these data provide the object autocorrelation function. If the object's Fourier phase can be retrieved, then the Fourier representation of the object will be complete and a 3-D Fast Fourier Transform (FFT) could be used to recover the object's 3-D complex reflectivity. Thus a phase-retrieval algorithm is an integral part of the PROCLAIM imaging modality.

We previously reported the development of an iterative phase-retrieval algorithm that relies on an opacity constraint. The algorithm seeks the maximum-likelihood estimate of the object's 3-D reflectivity, given measurements of the object's Fourier intensity that are corrupted with readout noise. Details of this algorithm are given in [4].

The performance of a phase-retrieval algorithm can be significantly enhanced by using good initial estimates. Such initial estimates can be found with the aid of *locator sets.* Locator sets are bounds on the object support derived from the support of the object's autocorrelation function. Crimmins, Fienup, Thelen, and Holsztynski first introduced the concept of locator sets along with various locator-set algorithms for use in performing phase retrieval from the autocorrelation of a two-dimensional object [7,8]. The generalization of these algorithms to three dimensions was subsequently demonstrated [5]. More mature locator-set algorithms for use with PROCLAIM data are described in a companion paper in this volume [9].

Locator sets can have a second function beyond providing initial estimates for phase retrieval. When they are tight enough, locator sets may be used directly for discrimination, identification, and even surface metrology. One can envision an alternative mode of operation for which the PROCLAIM acronym would more properly represent *Profile Retrieval* in place of *Phase Retrieval.*

The PROCLAIM sensor reflects the philosophy of reducing hardware complexity and costs at the expense of challenging algorithm development and increased computer processing. We believe that this is a favorable tradeoff, given the cost and performance trends in computer technology. As a consequence of this philosophy, there are no imaging optics, interferometric or heterodyne detection is avoided, precise mechanical and phase stability are not required, and laser coherence requirements are relaxed. A schematic diagram of the data-collection and processing that constitute the PROCLAIM imaging modality is presented in Figure 1.

3. PROCLAIM Simulation

In this section we present a simulation demonstration of the PROCLAIM imaging modality. For an object we use a simple cone that is illuminated perpendicular to its axis. Because this is an opaque object, it is completely characterized by its height and complex-reflectivity values as a function of cross-range location. Because it is represented by an array of numbers, the cone is discretely sampled in cross range, although the complex reflectivity can assume continuous values. The height has either discrete or continuous values depending on whether we use a voxel or a height-function representation. The cone has a half angle of 11.3 degrees, corresponding to an aspect that allows the height function to be conveniently stored in an real array of size 64×32 . Similarly, the complex reflectivity is stored in a 64×32 complex array. The complex reflectivity values were drawn from a circular complex Gaussian distribution, corresponding to a surface with diffuse scattering.

Data were simulated by taking the magnitude squared of the 3-D discrete Fourier transform of this cone object [4]. The dimensions of this data cube were $128 \times 64 \times 32$ corresponding to

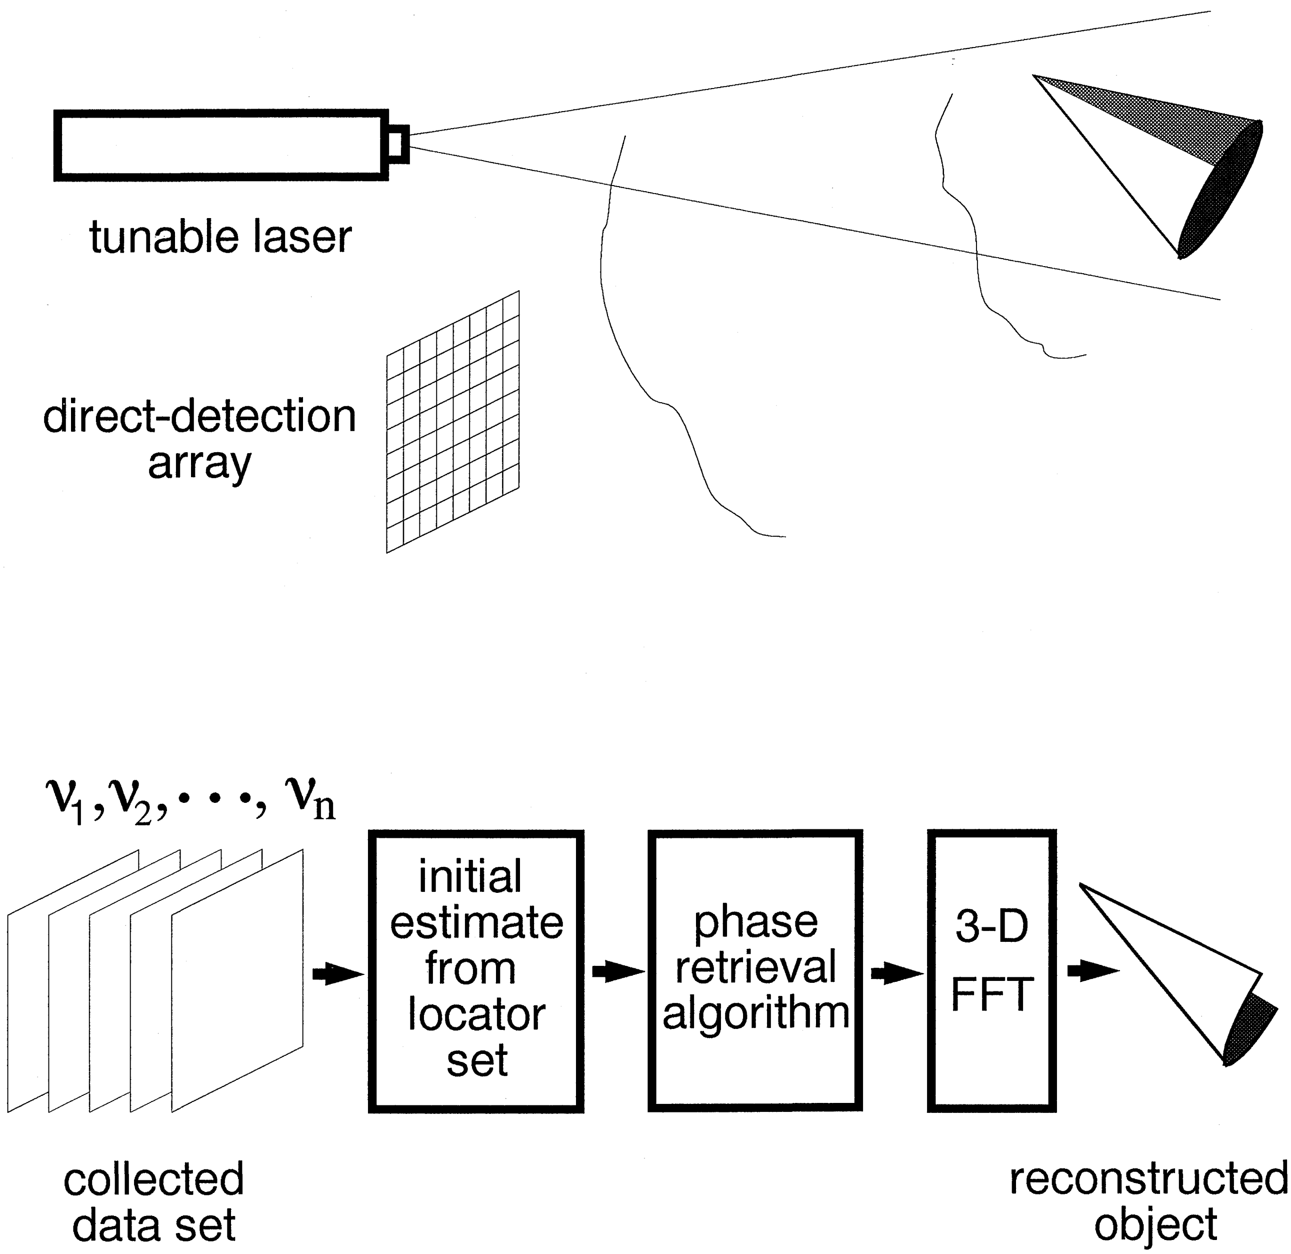

Figure 1: Data collection and processing sequence for the PROCLAIM sensor. The object is illuminated with a frequency-tunable laser and the far-field speckle intensity pattern is direct detected for each of multiple laser frequencies. Only the illuminated portion of the opaque object is reconstructed.

the extent of the autocorrelation of the cone. This data set is consistent with Nyquist-sampled far-field intensity data from a 128×64 focal-plane array, for each of 32 separate laser frequencies. For this demonstration, diffraction effects were not modeled and the data were uncorrupted by noise.

Our goal is to estimate a height value and a complex reflectivity value for each pixel in a 64×32 cross-range array. Our estimation algorithm is iterative and requires an initial estimate. Locator sets provide guidance in selecting an initial estimate. A locator set can be found by computing the object's autocorrelation function, thresholding to find the autocorrelation support, and applying locator-set algorithms to the autocorrelation support. Note that locator sets will be defined in terms of voxels. In our experience, use of mature locator-set algorithms for simulated data from cone objects of arbitrary aspect can produce locator sets that exactly match the opaque object support within the tolerance of a voxel representation. In this demonstration, however, we used a less accurate locator set that had residual thickness (i.e., multiple voxels in the range dimension). For the initial height estimate at each cross-range pixel, we used the midpoint of the locator set thickness. For the initial complex reflectivity estimate at each cross-range pixel, we used random circular complex Gaussian values (with a different seed than the true complex-reflectivity values).

The iterative estimation algorithm utilizes a conventional nonlinear optimization routine that maximizes the log-likelihood function over the height and complex reflectivity (real and imaginary parts) parameters at each cross-range pixel. At this point, the height is represented as a real number rather than by voxels. The search takes place in a parameter space of dimension 2130, corresponding to 710 cross-range pixels (defined by the locator set) each having a height parameter, and two complex reflectivity parameters. Initial estimates, final estimates, and true parameter values are shown in Figure 2. By comparing the individual pixels in (b) and (c) of Figure 2, one can see that the magnitude of the complex reflectivity is recovered well. A 3-D incoherent image could be reconstructed by incoherently averaging multiple reconstructions from disjoint data sets. By comparing the complex values (not shown), we find that the phase of the reconstruction is well matched to the true values, up to a single phase constant. High-fidelity estimates of the complex reflectivity are important because the complex reflectivity and height parameters interact: the fidelity of the complex reflectivity estimates helps to ensure the fidelity of the height estimates.

Initial and final height estimates along with the true height values can also be compared. In many applications it is the height function (profile information) that is of primary interest. The final estimate, Figure 2(e), compares favorably with the true height in (f), with an rms error of 0.054 pixels (range resolution elements) compared with 0.291 for the initial height estimate. It is apparent that the forced discretization in the initial estimate in (d) (resulting from the use of locator sets) has disappeared in the final estimate. This reconstruction demonstrates the potential for PROCLAIM to deliver superresolution in range.

4. PROCLAIM Laboratory Data Collection

The simulation described in the previous section demonstrates the feasibility of PROCLAIM. However, there is a significant difference between a successful simulation result and the demonstration of a technology on actual collected data. Measured data have a number of image-degrading effects, such as diffraction, speckle, focal-plane pattern noise, signal variation, read noise, bistatic

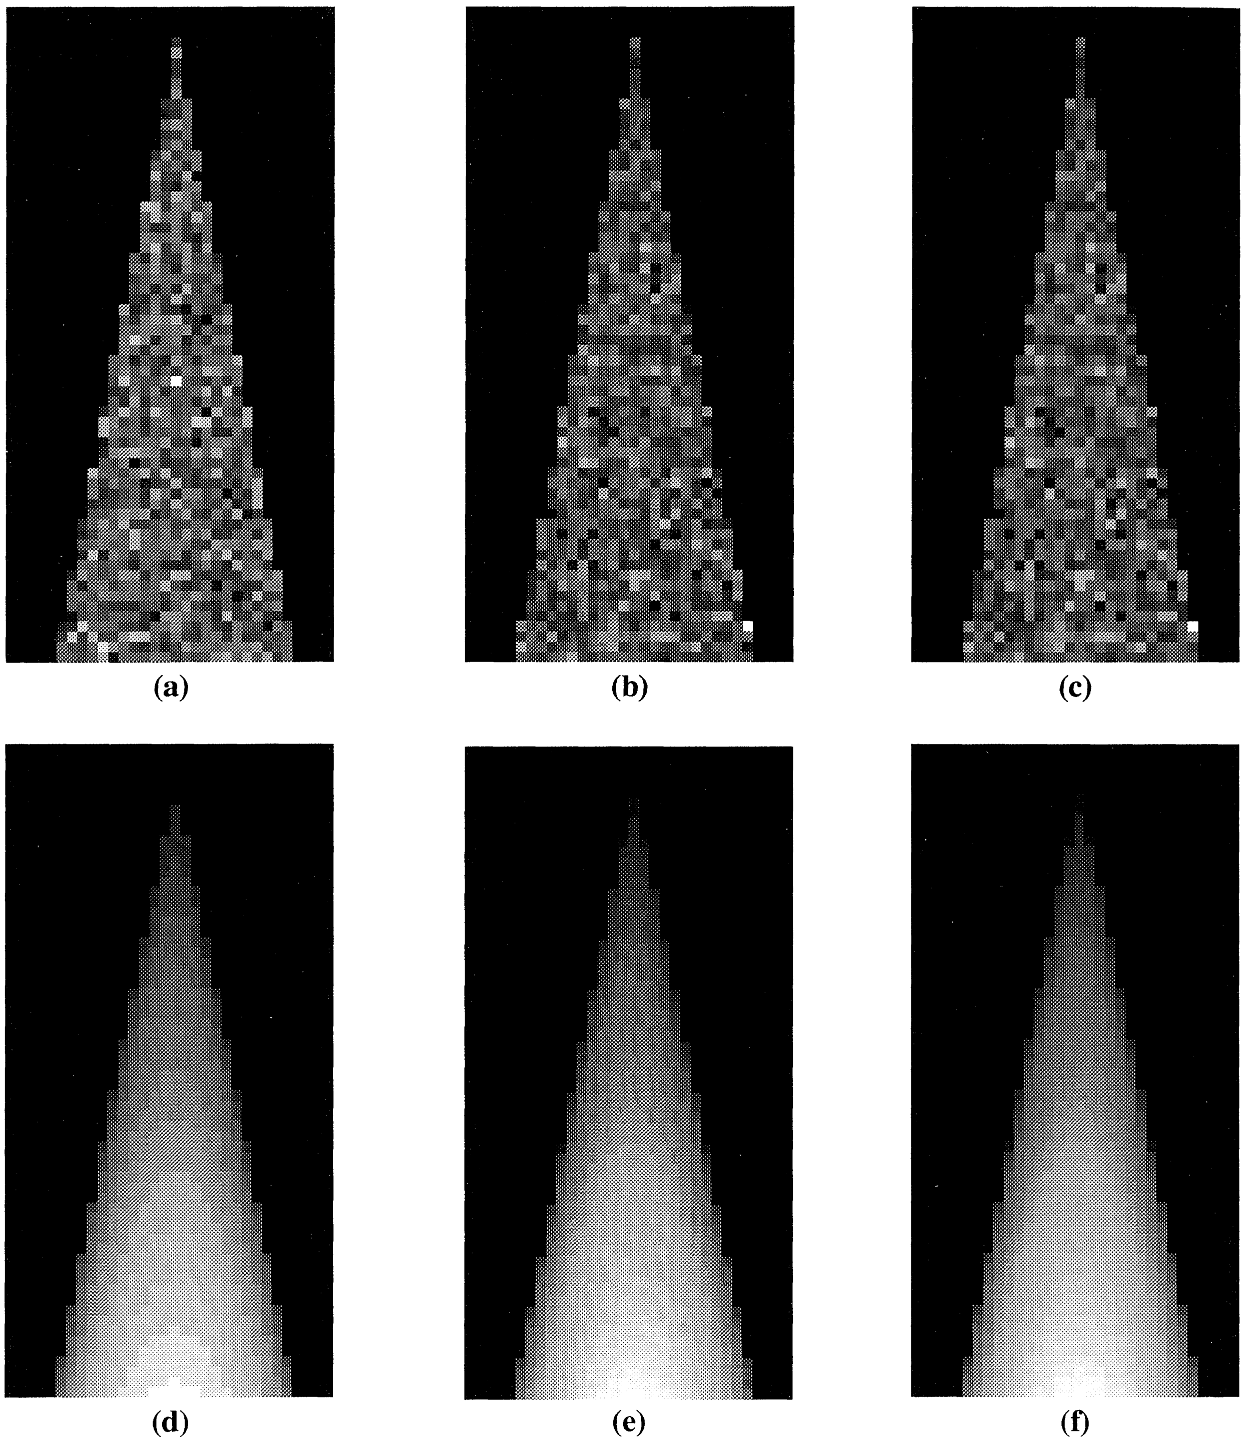

Figure 2 PROCLAIM reconstruction of cone object from simulated data. (a) Magnitude of initial estimate of complex reflectivity; (b) magnitude of reconstructed complex reflectivity; (c) magnitude of true complex reflectivity; (d) initial estimate of height, derived from locator set (brightness indicates height); (e) reconstructed height; (f) true height. The accuracy of the estimate for the complex reflectivity magnitude can be observed by comparing individual pixels in (b) and (c). Also, note that estimated height in (e) has continuous values that closely match the true height values in (f).

viewing geometry, and target reflectivity effects that are difficult to properly incorporate in a simulation model. Successful data reconstructions require that these potentially deleterious effects be overcome. In this section we describe the experimental setup, important data-collection considerations, and calibration and preprocessing steps that we found to be necessary to overcome the deleterious effects and to obtain accurate locator sets.

Our experimental setup consists of a near-infrared tunable diode laser, a 1024 × 1024, 12-bit CCD camera with 12μm detector pitch, and a target. The laser transmits in a stepped frequency mode. The beam flood-illuminates the target and the incoherent light is backscattered toward the detector array. A turning mirror is placed in the return path, close to the transmit beam to minimize the bistatic angle. A monostatic geometry using a beamsplitter could also be used.

Target size and range must be carefully chosen in accordance with the illuminating frequencies and the detector pitch. The maximum range extent of the target, T_r, must be less than one half of the range ambiguity interval because the autocorrelation of the target must fit completely within that interval:

$$T_r < \frac{\lambda^2}{4\Delta\lambda}, \qquad (1)$$

where $\Delta\lambda$ is the wavelength step.

The maximum angular extent of the target, T_a, is dictated by the minimum illumination wavelength, λ_{min}, the range to the detector, R, and the detector pitch, D_{pitch}:

$$T_a < \frac{R\lambda_{min}}{D_{pitch}}. \qquad (2)$$

We use a thin retroreflective tape to cover most of our targets. We find that this tape gives a strong diffuse return over large enough angles that the main lobe still falls on the detector for small bistatic angles.

We typically collect between 32 and 256 different wavelengths of data at wavelength steps of between 0.02 and 0.3 nm. More wavelengths correspond to a greater spatial-frequency bandwidth, and therefore finer range resolution.

In addition to the speckle data, we collect two frames of data of a flat background at two levels of incoherent illumination. We use these data to perform a two-parameter calibration to remove the fixed pattern noise due to pixel-to-pixel deviations in responsivity. This calibration eliminates significant noise found in the DC plane of the 3D autocorrelation.

A companion paper in this volume [9] describes in detail the PROCLAIM data processing steps and the intersection rules for computing object support bounds or locator sets. Briefly, we compute a 3-D inverse Fourier transform to generate the complex autocorrelation function, apply Spatially

Variant Apodization (SVA) to eliminate diffraction sidelobes, and incoherently average over multiple subapertures to reduce speckle nulls. Accurately determining the autocorrelation support is crucial to obtaining tight object support bounds.

In Figure 3 we show what the raw data looks like and the benefits of preprocessing for removing artifacts both outside and inside of the autocorrelation support. A portion of a raw data frame collected at a single frequency is shown in (a). A 128 × 128 subaperture is shown. In (b), we show a projection (i.e., the range dimension has been integrated out) of the 3-D autocorrelation of the object formed by performing a 3-D FFT on the 128 × 128 cross-range subaperture and all 64 frequencies. Note the diffraction sidelobes and the granularity of the support. In (c), the diffraction sidelobes are successfully removed by applying Spatially Variant Apodization (SVA). Finally, in (d), much of the granularity of the object support has been removed by incoherently averaging over 64 subapertures.

In the following section, we show results from applying our locator set algorithms to the preprocessed data sets.

5. PROCLAIM Locator Set Results

In this section, we apply our locator set rules to the 3-D autocorrelation support. These rules involve intersecting translated versions of the autocorrelation supports described in a companion paper [9]. The key to retrieving the profile support of an object from its autocorrelation is accurately determining the translation points. These points must be extreme points, meaning that there are no points in the support that lie beyond a plane that contains the candidate point. Additional intersections can be performed over z-thin points (i.e., where the autocorrelation support has just a single range value) if the cross-range support is known *a priori.*

Because the data are noisy, certain points can satisfy the extreme-point requirement, and yet not be part of the true autocorrelation support. For each candidate point we apply the following test to assess whether the point is part of the true autocorrelation support. We compute the intersection at the point under consideration (i.e., we shift the autocorrelation support center to that point of the current locator set and intersect) to form a new locator set, and then calculate the autocorrelation of the resulting new locator set. We expect that if the point belongs to the autocorrelation support, then the autocorrelation of this new locator set would be a superset of the measured autocorrelation support. Due to noisy data, this is not strictly true in practice, but we do expect that valid intersection points will not substantially violate this criterion. If the criterion is violated, then we do not use that point in for the new locator set. By applying the z-thin criterion, the extreme-point requirement, and the autocorrelation-support criterion, we find good points at which to perform intersections.

For one example, we choose an object which has a large cross-range extent relative to its surface relief. This is to demonstrate a fine range resolution appropriate for industrial inspection applications. The second example uses a conical object to demonstrate applicability of PROCLAIM to remote target characterization. The first object resembles an arrowhead (it is the difference between two isoceles triangles that have a common base, but different heights). It is relatively easy to find intersection points on the autocorrelation support of this shape. Computing a triple intersection at

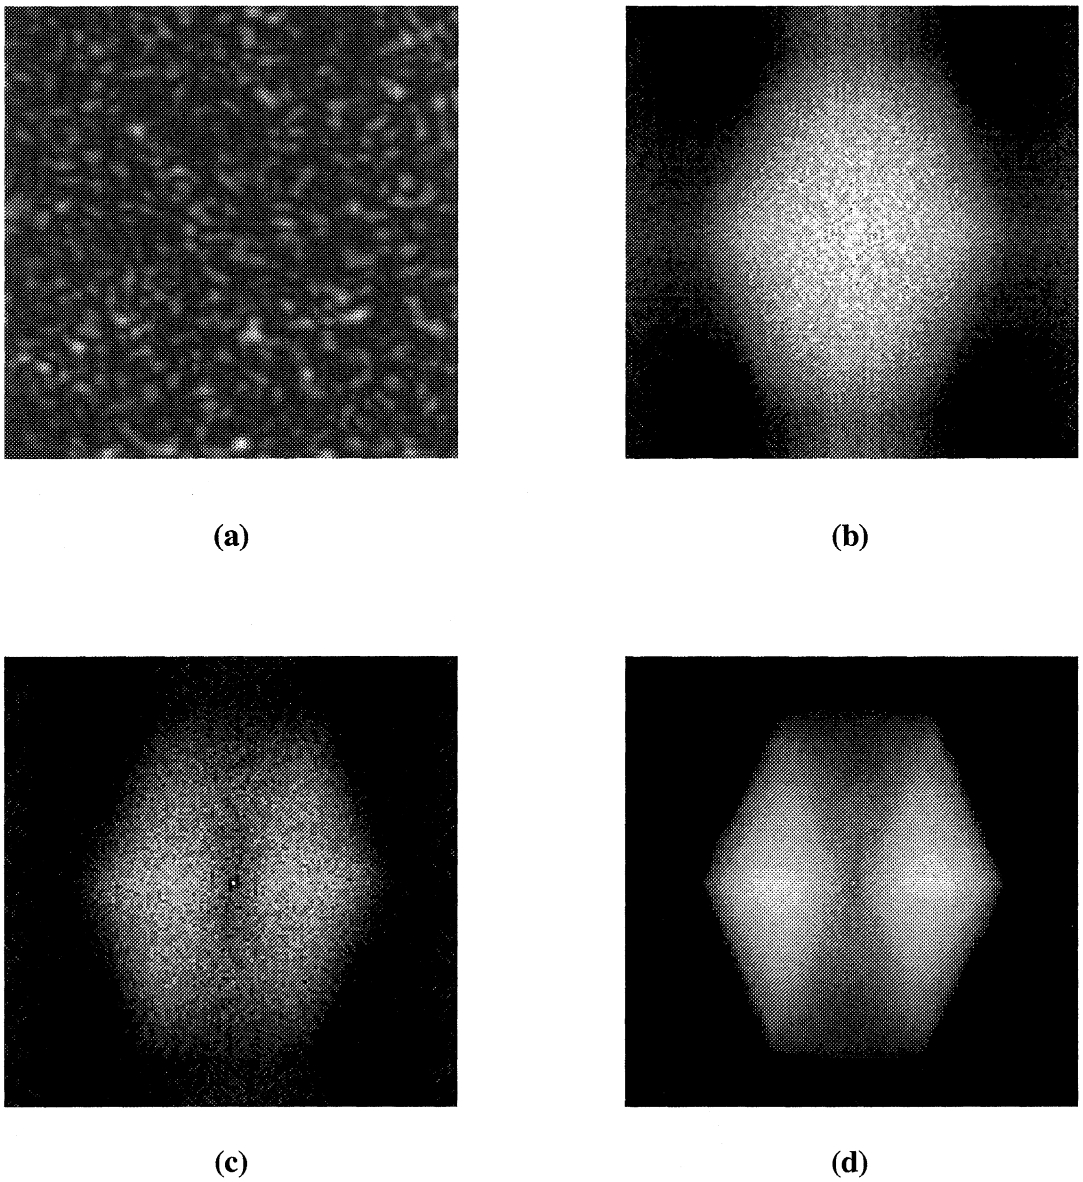

Figure 3 PROCLAIM laboratory data and preprocessing steps.(a) Subset of a raw data frame at a single frequency corresponding to the Fourier intensity of the backscattered signal; (b) projection of the 3-D autocorrelation of the data-cube generated with a 3-D FFT (brightness indicates integrated intensity of the autocorrelation function at a given cross-range pixel); (c) diffraction sidelobes removed by SVA; (d) incoherent average over multiple subapertures removes speckle nulls.

two of the extreme points readily yields an object support which closely resembles the object. We afixed three concentric circles on the arrowhead cut out of retroreflective tape. The thickness of the tape is 125 μm. We also afixed a retroreflective tape triangle on the right leg of the arrowhead to distiguish it from the other leg. Figure 4 (a) and (b) show the autocorrelation support and the locator set for this target. The concentric circles and the triangle are clearly discernable.

In the second example, we cover a cone with retroreflective tape. This object is less amenable to determining intersection points, but following our locator set rules using z-thin points, we are able to intersect over many points along the edge of the autocorrelation which corresponds to the base of the cone. Figure 4 (c) and (d) show the upper surface of the autocorrelation support and the locator set for the cone, respectively. The reconstructions are magnified by a factor of two relative to the autocorrelation supports to accentuate detail.

6. Conclusions

We have demonstrated, through simulation and with real data, a novel imaging modality called PROCLAIM. Intensity-only measurements are collected at multiple wavelengths using a laser and a lensless detector. Opacity is used to obtain a tight object support bound, which may be used as an initial estimate for a phase retrieval algorithm. PROCLAIM is applicable to problems ranging from industrial inspection to remote target characterization.

Acknowledgments

Supported by Ballistic Missile Defense Organization/ Innovative Science and Technology and managed by Avionics Directorate of Wright Laboratory, Aeronautical Systems Center, USAF, Wright Patterson AFB OH 45433.

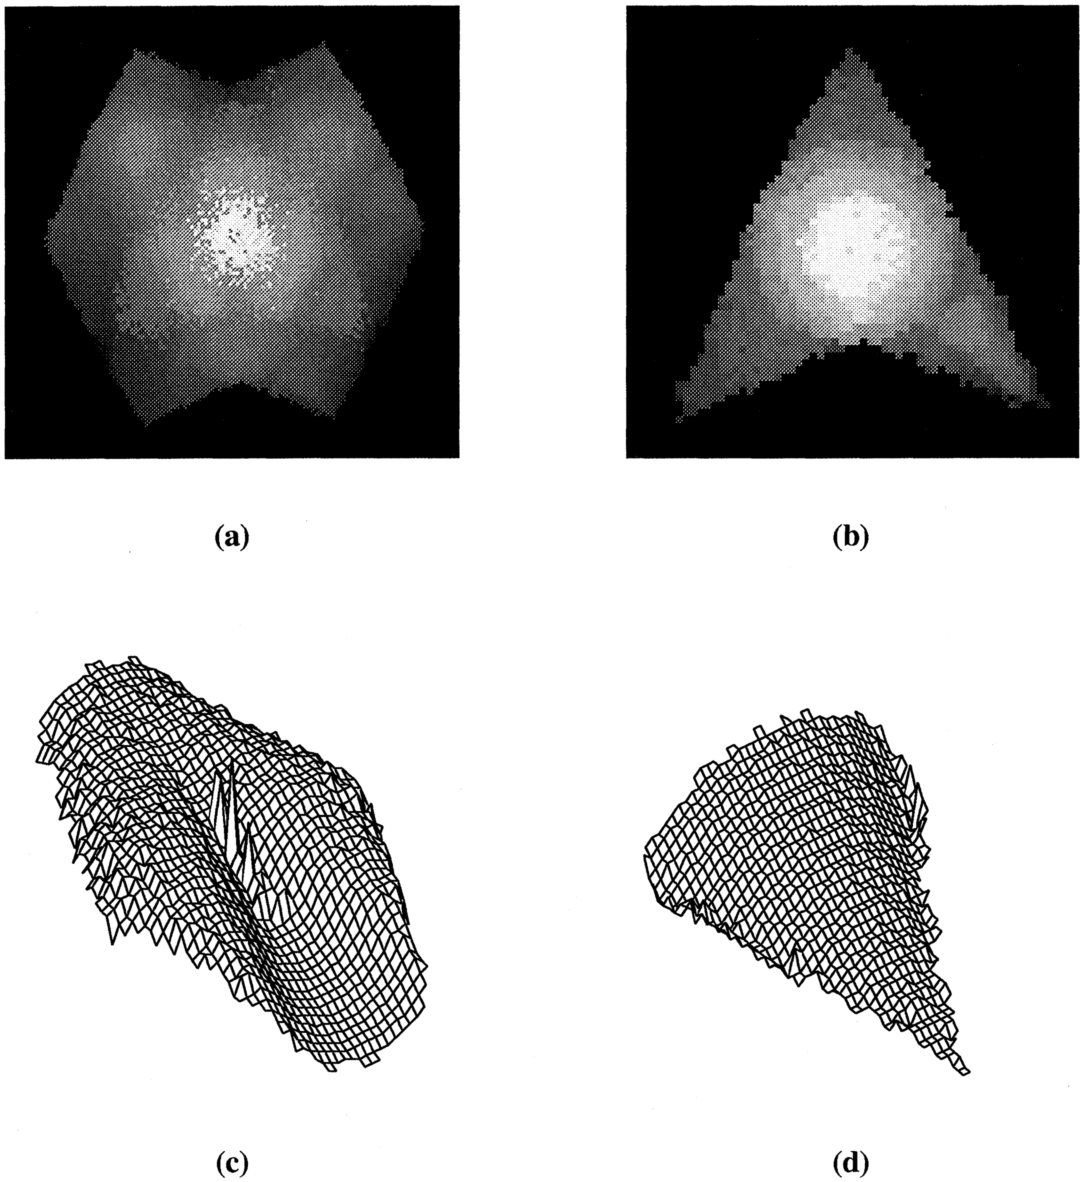

Figure 4: Autocorrelation supports and locator set results for experimental objects. (a) Upper surface of the autocorrelation support for arrowhead target with concentric circles (brightness indicates height); (b) resulting locator set formed by triple intersection; (c) upper surface of the autocorrelation support for cone target (perspective mesh plot); (d) resulting locator set formed by intersecting over multiple points along one edge of the autocorrelation support.

REFERENCES

1. J.H. Seldin, M.F. Reiley, R.G. Paxman, and B.E. Stribling, "Space-object identification using phase-diverse speckle," Invited paper in *Image Reconstruction and Restoration II* Proc. SPIE, (1997).
2. M.E. Bair, R. Sampson, and D. Zuk, "Three-dimensional imaging and applications," in *Intelligent Robots and Computer Vision*, D.P. Cassasent, ed., Proc. SPIE **726**, 264-274 (1987).
3. R.G. Paxman, "Superresolution with an opacity constraint," *Topical Meeting on Signal Recovery and Synthesis III*, Technical Digest Series **15**, (Optical Society of America, Washington DC, 1989), PD1, North Falmouth, Cape Cod, MA, June 1989.
4. R.G. Paxman, J.H. Seldin, J.R. Fienup, and J.C. Marron, "Use of an opacity constraint in three-dimensional imaging," in *Inverse Optics III*, Proc. SPIE **2241**, (1994).
5. R.G. Paxman, J.H. Seldin, J.R. Fienup, and J.C. Marron, "Phase retrieval with an opacity constraint," in *Topical Meeting on Signal Recovery and Synthesis V*, (Optical Society of America), 1995.
6. J.C. Marron and K.S. Schroeder, "Three-dimensional lensless imaging using laser frequency diversity," Appl. Opt. **31**, 255-262 (1992).
7. J.R. Fienup, T.R. Crimmins, and W. Holsztynski, "Reconstruction of the support of an object from the support of its autocorrelation," J. Opt. Soc. Am. A **72**, 610-624 (1982).
8. T.R. Crimmins, J.R. Fienup, and B.J. Thelen, "Improved bounds on object support from autocorrelation support and application to phase retrieval," J. Opt. Soc. Am. A **7**, 3-13 (1990).
9. J.R. Fienup, B.J. Thelen, M.F. Reiley, and R.G. Paxman, "3-D locator sets of opaque objects for phase retrieval," in *Image Reconstruction and Restoration II* Proc. SPIE 3170-10, (1997).