Closed-form Inverses for the Mixed Pixel/Multipath Interference Problem in AMCW Lidar

John P. Godbaz, Michael J. Cree and Adrian A. Dorrington

School of Engineering, University of Waikato, Hamilton, New Zealand

ABSTRACT

We present two new closed-form methods for mixed pixel/multipath interference separation in AMCW lidar systems. The mixed pixel/multipath interference problem arises from the violation of a standard range-imaging assumption that each pixel integrates over only a single, discrete backscattering source. While a numerical inversion method has previously been proposed, no close-form inverses have previously been posited. The first new method models reflectivity as a Cauchy distribution over range and uses four measurements at different modulation frequencies to determine the amplitude, phase and reflectivity distribution of up to two component returns within each pixel. The second new method uses attenuation ratios to determine the amplitude and phase of up to two component returns within each pixel. The methods are tested on both simulated and real data and shown to produce a significant improvement in overall error. While this paper focusses on the AMCW mixed pixel/multipath interference problem, the algorithms contained herein have applicability to the reconstruction of a sparse one dimensional signal from an extremely limited number of discrete samples of its Fourier transform.

Keywords: Mixed Pixels, Multipath Interference, AMCW, Lidar, Attenuation Ratio, Range Imaging, Systematic Error, Multiple Return Separation

1. INTRODUCTION

Full-field Amplitude Modulated Continuous Wave (AMCW) lidar is an important new development in ranging technology. AMCW lidar operates on the time-of-flight principle: by illuminating a scene with modulated lasers or LEDs and measuring the delay introduced in the modulation envelope by varying target distances, it is possible to produce high precision range measurements that can be converted to three dimensional point-clouds. Whereas previous systems required a physical mechanism to explicitly scan a point across a scene, advanced new systems such as the SwissRanger SR4000[1] and Canesta XZ-422 produce 2D matrices of range data in a single, simultaneous measurement process. These 'range images' amount to instant 3D models of a target or scene and could potentially be applied to such divergent applications as face recognition, computer gaming, robotics, special effects production for movies and manufacturing process quality control.

However, full-field ranging is still a developing technology and there are a number of significant technical problems to be overcome. In this paper we discuss two new approaches to the mixed pixel/multipath interference problem.[2] The simultaneous imaging of a large number of points drastically increases the problem of measurement cross-talk, primarily caused by scattering within the imaging optics of full-field rangers.[3,4] Similar problems occur due to the simultaneous illumination of the entire scene, whereby multiple reflections within the scene potentially result in multipath interference,[5] resulting in erroneous range measurements. Lastly, the imaging nature of full-field systems increases the impact of mixed pixels, caused when a single pixel integrates over a region at the boundary of an object and contains two or more target objects at different ranges.[6] Despite many commercial systems advertising sub-centimetre precision, in practice the accuracy can be one or more orders of magnitude worse due to systematic errors, primarily due to multipath interference.

Mixed pixels have been long identified as a problem in point scanning systems,[6] however research has primarily focussed on detection and removal of perturbed points, rather than correction of points. More recent work on full-field systems has applied diverse approaches to multipath reduction: several authors have assumed relatively spatially homogeneous scattered light across the entire image and removed the impact of the scattered light by

---

Further author information: J.P.G. Email: jpg7@waikato.ac.nz, Postal Address: School of Engineering, University of Waikato, Private Bag. 3105, Hamilton, 3240, New Zealand

using calibration squares within the scene or by detecting how scene texture/patterning or structured illumination affects the detected ranges to objects.[7, 8] For example, for an object which is half black and half white, the black region will appear more highly perturbed by any scattered light then the white region. Other authors have attempted to model the cross-talk between measurements introduced by intra-camera scattering,[3, 9] however this problem is highly spatially variant, which makes accurate inversion an extremely difficult deconvolution problem. Fuchs[10] introduced a particularly interesting approach to ameliorating multipath due to intra-scene scattering: by explicitly modelling the reflection of the modulated illumination within the scene, he estimated the perturbations introduced by multiple reflections and removed them. However, this type of approach is limited in that it is only possible to model the effect of scattering within the field-of-view, modelling occlusion is highly computationally complex and the specularity of objects is unknown.

Whereas most previous amelioration methods have operated on either mixed pixels or multipath interference, but not both, there are two other approaches which operate on both. The first method uses the harmonic content of the AMCW correlation waveform to separate out different range sources or component returns within each measurement;[11, 12] if the darker components are likely to be cross-talk from other pixels, then a 'correct' unperturbed range value can be calculated. The second method uses measurements at two different modulation frequencies to allow up to two different contributing component returns to be discriminated,[13] which has a distinct similarity to the methods presented in this paper except that the authors did not provide a closed-form solution, instead relying on numerical optimisation to separate out the different backscattering component returns. Relying on numerical optimisation makes analysis of the noise properties of the solution more difficult in addition to dramatically increasing the computational complexity of the algorithm. A truly real-time solution requires a closed-form solution.

While this paper focusses on the AMCW mixed pixel/multipath interference problem, the algorithms contained herein have applicability to the reconstruction of any sparse one dimensional signal from an extremely limited number of discrete samples of its Fourier transform.

In the next section we address the formation of AMCW lidar range measurements and demonstrate how multiple backscattering sources within a pixel results in perturbed range and amplitude estimates. Section 3 introduces a model for distributed reflectivity/scattering, such as occurs when making measurements of scenes subject to fog. By careful choice of model it is possible to derive a set of polynomial equations which can be simultaneously solved in closed form, resulting in a new Cauchy distribution based separation method. Section 4 introduces attenuation ratios and derives a general class of polynomial equations based on these ratios. Simultaneous solution of these attenuation ratio polynomials gives a second closed-form mixed pixel/multipath interference separation algorithm.

2. BACKGROUND

This section models the formation of AMCW range measurements. For the sake of simplicity, a minimalist model is presented: for more detailed information on the AMCW lidar measurement formation process see[2] and.[14]

We utilise the following conventions: the Fourier transform of a function $f(x)$ is notated as $F(u)$, square brackets, $[\,]$, are used to notate functions of a discrete variable, $*$ is the complex conjugate and $j^2 = -1$.

2.1 Modelling Measurements

The focus of this paper is the recovery of a function $f_\xi(r) \in \mathbb{R}^+$, representing the lidar signal returns within a single pixel, from N samples of the Fourier transform, $F_\xi(u)$. A similar problem is addressed in full-waveform lidar, where a similar model of backscattered returns is recovered; however, rather than measuring a convolution of f_ξ with an impulse response, we only have a small number of measurements of discrete spatial frequencies of the signal returns.

While it is relatively easy to find a closed-form inverse that estimates a discrete model of the function $f_\xi(r)$, it is much more difficult to find an inverse that produces continuous range information. As the number of component returns/backscattering sources increases, the problem becomes substantially more difficult to solve in closed-form. As a result, we concentrate on the case of two components and N measurements.

Restating our measurement formation model from,[2, 14] $f_\xi(r)$ can be considered to be signal intensity as a function of range for a single pixel. Let the illumination modulation waveform be $\psi_i(\phi)$, a function of relative phase. As light travels from the camera to the scene and back it encodes distance travelled as a phase shift in the illumination. If there is more than one backscattering source within a pixel, then the illumination waveform measured at the sensor is a superposition of the illumination reflected by each of the backscattering sources, hence the illumination measured at the sensor can be modelled as a convolution of the illumination modulation waveform with the backscattering intensity as a function of range. This gives

$$\Psi_m(u) = F_\xi\left(\frac{4\pi}{\lambda}u\right)\Psi_i(u), \tag{1}$$

where Ψ_m is the Fourier transform of the backscattered illumination at the sensor and λ is the wavelength of the illumination modulation. This backscattered illumination is indirectly measured by correlating it with a reference signal, $\psi_s(\phi)$, giving the correlation waveform defined by

$$H(u) = \Psi_m(u)\Psi_s^*(u), \tag{2}$$

where $H(u)$ is the Fourier transform of the correlation waveform. By measuring the correlation waveform while changing the relative phase of the reference signal, it is possible to measure a discrete spatial frequency of $f_\xi(r)$ and determine range and amplitude. In an ideal case an AMCW lidar measurement, $\xi \in \mathbb{C}$, is formed by sampling the fundamental frequency of the correlation waveform and correcting for the influence of modulation waveform shape, namely

$$\xi = \frac{H(-1/2\pi)}{\Psi_i(-1/2\pi)\Psi_s^*(-1/2\pi)} = F_\xi\left(-\frac{\lambda}{2}\right), \tag{3}$$

which is equivalent to sampling the $-2/\lambda$ spatial frequency of the backscattering function, $f_\xi(r)$. The standard assumption for an AMCW lidar imaging system is that there is only ever a single component return within each pixel, with amplitude a_0 and range d_0. In this case the backscattering function is expressed by

$$f_\xi(r) = a_0\delta(r - d_0), \tag{4}$$

subsequently, a calibrated AMCW measurement at a wavelength of λ is given by

$$\xi = F_\xi(-2/\lambda) = a_0 e^{4\pi j d_0/\lambda}, \tag{5}$$

and the amplitude of, a_0, and range to, d_0, the object imaged by the pixel can be recovered from a complex domain measurement by

$$d_0 = \frac{\lambda}{4\pi}(\arg(\xi) + 2\pi m) \tag{6}$$

$$a_0 = |\xi|, \tag{7}$$

where $m \in \mathbb{Z}$ is a phase unwrapping constant. However, this method of range and amplitude recovery breaks down in the presence of multiple component returns: these extra components are the source of the mixed pixel and multipath interference problems.

2.2 The Mixed Pixel/Multipath Interference Case

In the simplest mixed pixel/multipath interference case, there are two backscattering sources, giving a measurement that is the sum of contributions from two components,

$$\xi = a_0 e^{4\pi j d_0/\lambda} + a_1 e^{4\pi j d_1/\lambda}. \tag{8}$$

We primarily focus on this case in this paper. The perturbation introduced in the AMCW measurement by the second component can be modelled by a function $\Lambda_f(b, \theta)$, defined as

$$\Lambda_f(b, \theta) = \frac{\xi}{a_0 e^{4\pi j d_0/\lambda}} = 1 + be^\theta, \tag{9}$$

where $b = a_1/a_0$ is the relative amplitude of the darker return (where $a_0 \geq a_1$) and $\theta = 4\pi(d_1 - d_0)/\lambda$ is the relative phase.

In order to deal with multiple component returns, it is necessary to develop a more advanced model for $f_\xi(r)$ and a manner to fit it to measured data. One way to address this is by taking multiple measurements at different modulation frequencies. There are two ways to achieve this: one is to take explicit measurements at different modulation frequencies,[13] the other method is to implicitly take measurements at different modulation frequencies using the harmonic content of the correlation waveform.[11] Previous work has applied multiple measurements at different modulation frequencies and the synthetic wavelength technique to the unwrapping of range measurements;[15] until recently, this was the only manner in which multiple measurements at different modulation frequencies were utilised.

One simple way to recover $f_\xi(r)$ from a sequence of measurements of the Fourier transform is to systematically sample at different modulation frequencies and then use an inverse Fourier transform to reconstruct a model of the signal returns. For example, Simpson *et al.*[16] sampled from 10 MHz to 200 MHz at 10 MHz intervals and reconstructed a model in this exact manner. However, this requires a large number of measurements across a very large bandwidth, which is very challenging to achieve with high efficiency and accurate calibration, and does not provide particularly high spatial resolution given the number of samples required. Another potential approach is to use an explicit prior distribution (for example, a Gaussian intensity distribution) to help fill-in-the-blanks and reconstruct a discrete model from an incomplete sequence of measurements at different modulation frequencies. While this is a standard signal-processing approach to reconstruction, it does not appear to have been previously applied to reconstruction of $f_\xi(r)$ in AMCW lidar.

Another approach is to pose the recovery of $f_\xi(r)$ as a sparse-spike-train deconvolution problem and apply an off-the-shelf deconvolution method.[11] In the case of,[11] the Levy-Fullagar algorithm, although there are many other possible approaches that could be applied such as $\mathcal{L}_1$ norm regularisation and single-most-likely replacement type methods. However, all these methods are complicated iterative processes – a closed-form inverse is preferable, in particular, a closed-form inverse that inherently determines continuous range, rather than a discrete model of signal returns such as that returned by deconvolution. While it is possible to determine continuous range values from a discrete model, it tends to be a difficult process fraught with assumptions.

3. INVERSESE MODELLING THE DISTRIBUTION OF REFLECTIVITY

For real scenes it is reasonable to assume that the backscattering function is sparse; reflective objects tend to occlude other objects behind them and in any interesting scene there is space between the camera and the object. On the other hand, it is naïve to assume that component returns are infinitesimal, such as assumed by eqn. 4. Imaging an object, other than a perfect sphere centred at the camera, always results in collecting light from a finite range of radial distances. In addition, there are also volumetric scattering sources such as fog, that result in distributed component returns. Combining these two assumptions leads to a general model for $f_\xi(r)$ of n discrete component returns as

$$f_\xi(r) = \sum_{i=0}^{n-1} a_i g_i(r - d_i), \tag{10}$$

where a_i is the amplitude of the ith component return, d_i is the range to the return and $g_i(r)$ is a model of how diffuse the reflectivity is. In this case, making assumptions about the structure of the backscattered signal returns allows a more accurate reconstruction.

Given N measurements, where measurement l is at modulation frequency $f_l = r_l\gamma$, where the relative frequencies $r_l \in \mathbb{Z}$ are coprime and γ is an arbitrary base frequency, a measurement ξ_l at a relative frequency r_l can be expressed in terms of component returns notated at a relative frequency of unity. The ith component return, $\eta_i \in \mathbb{C}$, is then

$$\eta_i = a_i \mathrm{e}^{4\pi d_i \gamma/c}, \tag{11}$$

where a_i and d_i are the amplitude and range to the return and c is the speed of light. This is the contribution to the complex domain range measurement by the ith return if measured at a relative frequency of unity. From basic

principles, if component return i is measured at relative frequency r_l, it will have an apparent phase $r_l \arg(\eta_i)$. Using eqn. 10, ξ_l can be notated as

$$\xi_l \;=\; F_\xi\left(-r_l\frac{\gamma}{2c}\right) \tag{12}$$

$$=\; \sum_{i=0}^{N-1} a_i G_i\left(-r_l\frac{\gamma}{2c}\right) \mathrm{e}^{r_l \arg(\eta_i)} \tag{13}$$

$$=\; \sum_{i=0}^{N-1} G_i\left(-r_l\frac{\gamma}{2c}\right) \frac{\eta_i^{r_l}}{|\eta_l|^{r_l-1}}. \tag{14}$$

Eqn. 14 ultimately leads to a class of simultaneous polynomials written in a Cartesian form, splitting the real and imaginary parts of η_i. Using the η_i notation avoids either the removal or addition of any unnecessary cyclic ambiguity; the number of possible solutions for $H = \{\eta_0, \eta_1, ..., \eta_{n-1}\}$ given a set of measurements, $\Xi = \{\xi_0, \xi_1, ..., \xi_{N-1}\}$, exactly corresponds to the number of physically possible solutions. Expressing each return in a Cartesian form,

$$\eta_i = R_i + jI_i, \tag{15}$$

where $R_i, I_i \in \mathbb{R}$, allows us to rewrite eqn.14 as

$$\xi_l = \sum_{i=0}^{n-1} \frac{(R_i + jI_i)^{r_l}}{(R_i^2 + I_i^2)^{(r_l-1)/2}}. \tag{16}$$

In the simplest possible case of a Dirac delta model ($g(r) = \delta(r)$ thus $G(u) = 1$), two measurements and two components, this equation can be converted to polynomial form by rearranging to remove the fractional powers, with higher relative frequencies resulting in increased polynomial order. The two return case for positive, odd values of the relative frequency, $r_l \in \mathbb{Z}$, expands to

$$\xi_l S_0^{(r_l-1)/2} S_1^{(r_l-1)/2} - U_0^{r_l} S_1^{(r_l-1)/2} - U_1^{r_l} S_0^{(r_l-1)/2} = 0, \tag{17}$$

where $U_i = R_i + jI_i$ and $S_i = R_i^2 + I_i^2$. This polynomial is r_lth order over each of R_0, R_1, I_0 and I_1, and if necessary can be written explicitly as separated real and imaginary equations. The same expansion for even relative frequencies results in a polynomial of $4r_l$th order, which is substantially worse. Apart from a ratio of 3:1, every single frequency ratio results in at least one polynomial of fifth order or more, even assuming that there was some nice way to impose the constraint that each component is real there is unlikely to be a general inverse. Writing out the equivalent polar form, however, is trivial. Rewriting eqn. 13 gives

$$\xi_l \;=\; \sum_{i=0}^{N-1} a_i G_i\left(-r_l\frac{\kappa}{2c}\right) \kappa_i^{r_l}, \tag{18}$$

where $a_i \kappa_i = \eta_i$, such that $|\kappa_i| = 1$ encodes the range information and a_i the amplitude. A relative frequency of r_l gives a polynomial with order of r_l over each κ_i. However, in the same way that there was no easy way to impose the constraint that R_i and I_i be real in the case of the Cartesian model, there is no easy way to impose the constraints that a_i be real and κ_i be constrained to the unit circle. Even though two measurements at different modulation frequencies is enough to exactly determine two component returns using optimisation, in order to solve the problem as simple simultaneous polynomials requires that the problem be effectively overdetermined.

3.1 A Simple Inverse (Dirac Delta Model, Four Measurements, Two Returns)

Taking four measurements at relative frequencies, $r_l \in \{\rho, \rho+1, \rho+2, \rho+3\}$, where $\rho \in \mathbb{Z}$ is an arbitrary constant, allows the problem in the Dirac delta model case to be notated as

$$\xi_0 \;=\; \mu_0 + \mu_1 \tag{19}$$

$$\xi_1 \;=\; \mu_0 \kappa_0 + \mu_1 \kappa_1 \tag{20}$$

$$\xi_2 \;=\; \mu_0 \kappa_0^2 + \mu_1 \kappa_1^2 \tag{21}$$

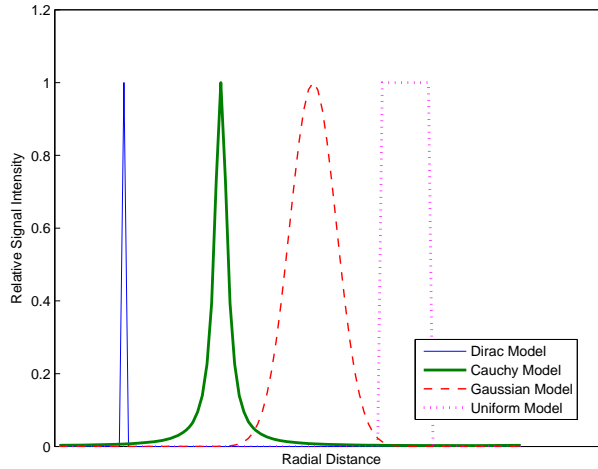$$\xi_3 \;=\; \mu_0 \kappa_0^3 + \mu_1 \kappa_1^3, \tag{22}$$

Figure 1: Candidate diffuse-reflectivity return models

where $\mu_i = a_i \kappa_i^\rho$ is an intermediate variable. From eqn. 19,

$$\mu_1 \;=\; \xi_0 - \mu_0. \tag{23}$$

$$\Rightarrow \quad a_1 \;=\; \frac{\xi_0 - \mu_0}{\kappa_1^\rho}. \tag{24}$$

Substituting eqn. 23 into eqn. 20 gives

$$\mu_0 \;=\; \frac{\xi_0 \kappa_1 - \xi_1}{\kappa_1 - \kappa_0}, \tag{25}$$

$$\Rightarrow \quad a_0 \;=\; \frac{\xi_0 \kappa_1 - \xi_1}{(\kappa_1 - \kappa_0)\kappa_0^\rho} \tag{26}$$

which combined with eqn. 21 allows the determination that

$$\kappa_1 = \frac{\xi_1 \kappa_0 - \xi_2}{\xi_0 \kappa_0 - \xi_1}. \tag{27}$$

Substituting these equations into eqn. 22 gives a polynomial equation, namely,

$$f\kappa_0^2 + g\kappa_0 + h = 0, \tag{28}$$

where $f = \xi_0 \xi_2 - \xi_1^2$, $g = \xi_1 \xi_2 - \xi_0 \xi_3$ and $h = \xi_1 \xi_3 - \xi_2^2$. This can be trivially solved to give two different solutions corresponding to the two different permutations of the same two component returns. This is a nice solution in that there is a direct inverse, but is problematic in a practical sense because it requires four separate measurements when two is enough to fully constrain the problem. One mitigating aspect, is that by measuring the extent to which the positive real number constraints on $\{R_0, R_1, I_0, I_1\}$ are broken, may give a measure of how much error there is in the coefficient estimates.

3.2 Models For Diffuse-Reflectivity

There are many different plausible models for diffuse-reflectivity, several examples are plotted in fig. 1. A Gaussian model is a very natural model for component returns that originate from a region of space. Gaussian models are commonly fit to measurements in pulsed radar and lidar[17] in order to determine the range and amplitude of component returns. These systems typically involve direct sampling in the spatial domain, whereas

AMCW lidar relies on sparse sampling directly in the Fourier domain and common numerical fitting approaches such as Expectation Maximisation[18] are not guaranteed to converge to the correct answer due to local minima.

The zero centred Gaussian distribution with a variance of σ^2 is given by

$$\mathcal{N}(0, \sigma^2) = \frac{1}{\sqrt{2}\sigma} e^{-\frac{x}{2\sigma^2}}, \tag{29}$$

with a Fourier transform of

$$\mathcal{F}\{\mathcal{N}(0, \sigma^2)\}(u) = \sqrt{\pi} e^{-2\pi^2 u^2 \sigma^2}. \tag{30}$$

The attenuation constant encodes the rate at which the specific Gaussian distribution attenuates modulation frequencies as modulation frequency increases. From the preceeding equation, we can write the attenuation constant in the Gaussian case as

$$\nu = e^{-2\pi^2 \sigma^2}, \tag{31}$$

enabling eqn. 18, assuming Gaussian distributed returns, to be expressed as

$$\xi_l = \sum_{i=0}^{N-1} a_i \nu_i^{r_l^2} \kappa_i^{r_l}. \tag{32}$$

However, like the Cartesian Dirac delta model, the order of the polynomial increases quickly as r_l increases. In the single component, three measurement case, there is an inverse given by

$$\nu_0 = \left(\frac{|\xi_2|}{|\xi_1|} \right)^{\frac{1}{r_2^2 - r_1^2}} \tag{33}$$

$$\kappa_0 = \left(\frac{\xi_1}{\xi_0 \nu_0^{r_1^2 - r_0^2}} \right)^{\frac{1}{r_1 - r_0}} \tag{34}$$

$$a_0 = \frac{\xi_0}{\nu_0^{r_0^2} \kappa_0^{r_0}}. \tag{35}$$

However, there is no simple closed form for more than one component, making mixed pixel/multipath interference separation difficult for a Gaussian model.

While many other distributions, such as a uniform distribution defined by a rectangular function of variable width, do not have exact polynomial expressions, the Cauchy distribution does. The Probability Distribution Function of a zero centred Cauchy-Lorentz distribution is given by

$$f_\mathcal{C}(x|0, \lambda_\mathcal{C}) = \frac{\lambda_\mathcal{C}}{\pi(x^2 + \lambda_\mathcal{C})}, \tag{36}$$

where $\lambda_\mathcal{C}$ is the width of the distribution. The Fourier transformation of a Cauchy distribution is given by a zero centred Laplace distribution, viz

$$\mathcal{F}\{f_\mathcal{C}(x|0, \lambda_\mathcal{C})\}(u) = f_\mathcal{L}(u|0, \lambda_\mathcal{C}) = \frac{1}{2\lambda_\mathcal{C}} e^{-\frac{|u|}{\lambda_\mathcal{C}}}. \tag{37}$$

Whereas it was previously assumed that $|\kappa_i| = 1$, we now utilise the modulus of κ_i to encode the attenutating properties of the Cauchy distribution in the Fourier domain. In other words, let $|\kappa_i|$ be given as

$$|\kappa_i| = e^{-1/\lambda_\mathcal{C}}, \tag{38}$$

thus the relationship between the modulii of measurements of a component return at relative frequencies r_0 and r_1 are given by

$$\frac{f_\mathcal{L}(r_1|0, \lambda_\mathcal{C})}{f_\mathcal{L}(r_0|0, \lambda_\mathcal{C})} = \frac{|\kappa_i|^{|r_1|}}{|\kappa_i|^{|r_0|}}, \tag{39}$$

allowing eqn. 18 to be expressed in the Cauchy case for positive r_l as

$$\xi_l \quad = \quad \sum_{i=0}^{N-1} a_i \kappa_i^{r_l}, \tag{40}$$

and for negative r_l as

$$\xi_l^* \quad = \quad \sum_{i=0}^{N-1} a_i \kappa_i^{|r_l|}. \tag{41}$$

It can be clearly seen that the Dirac delta model is merely a degenerate case of the Cauchy model, when $|\kappa_i| \to 1$. Taking into account the conjugate relationship for negative relative frequencies, the inversion formula from eqns. 23 to 28 is also applicable to the Cauchy model. Whereas for a Dirac delta model, three complex and one positive real domain measurements are required in order to determine four positive real values, in this case three complex and one positive real domain measurements are required in order to determine two real and two complex domain values. This is significantly less inherently overdetermined, thus more efficient; albeit, in practice all measurements must be treated as complex.

4. INVERSION VIA ATTENUATION RATIO POLYNOMIALS

In the previous section we introduced the idea of modelling range-diffuse measurements using simultaneous multivariate polynomial equations. While in a number of overdetermined cases there are closed-form solutions, generally the polynomial equations involved are too complicated and high order for a closed-form solution. This section introduces a different representation for Dirac delta type point-returns using the attenuation ratio. This allows the relationship between component returns to be analysed independently of the absolute phase and amplitude; as a result, the polynomial order can be reduced.

4.1 The Attenuation Ratio

Let the attenuation ratio τ_l at a particular modulation frequency vr_l be

$$\tau_l = \frac{|\xi_l|}{\omega} = \frac{\left| \sum_{i=0}^{n-1} a_0 \kappa^{r_l} \right|}{\sum_{i=0}^{n-1} a_0}, \tag{42}$$

where ω is the total backscattered intensity, equal to the sum of the modulii of the component returns, and assuming a Dirac delta model where $|\kappa_i| = 1$. Potentially, this can be approximated by taking the modulus of a measurement at an extremely low modulation frequency. In the two return case, assuming that

$$\omega \quad = \quad a_0 + a_1 \quad = \quad a_0(1+b), \tag{43}$$

the attenuation ratio can be modelled using eqn. 9 by

$$\tau_l = \frac{a_0 |\Lambda_{\mathrm{f}}(b, r_l\theta)|}{a_0(1+b)} = \frac{|\Lambda_{\mathrm{f}}(b, r_l\theta)|}{1+b}. \tag{44}$$

The attenuation ratio is the ratio of net amplitude to the sum of the component amplitudes and indicates how much cancellation has occured between component returns for a particular pixel. A geometric interpretation of eqn. 9 combined with the law of cosines gives

$$|\Lambda_{\mathrm{f}}(c, \theta)|^2 = 1 + c^2 + 2c\cos(\theta). \tag{45}$$

Assume two measurements, ξ_0 and ξ_1, at relative frequencies, r_0 and r_1, with known attenuation ratios, τ_0 and τ_1. Solving eqn. 44 to find $|\Lambda_{\mathrm{f}}(b, \theta)|$ and combining with eqn. 45 gives

$$|\Lambda_{\mathrm{f}}(b, r_l\theta)|^2 = 1 + b^2 + 2b\cos(r_l\theta) = \tau_l^2(1+b)^2, \tag{46}$$

hence

$$b^2 + b\frac{2\cos(r_l\theta) - 2\tau_l^2}{1 - \tau_l^2} + 1 = 0. \tag{47}$$

This is the attenuation ratio polynomial for a measurement at a particular relative frequency r_l. Given a sequence of attenuation ratios at different spatial frequencies, a set of polynomials is formed which can be simultaneously solved; we now address the solution of these polynomial systems.

4.2 Determining Relative Phase

Eqn. 47 is separately valid for each measurement, thus

$$\frac{2\cos(r_0\theta) - 2\tau_0^2}{1 - \tau_0^2} = \frac{2\cos(r_1\theta) - 2\tau_1^2}{1 - \tau_1^2}, \tag{48}$$

or

$$\alpha\cos(r_1\theta) - \beta\cos(r_0\theta) + \gamma = 0, \tag{49}$$

where $\alpha = 1 - \tau_0^2$, $\beta = 1 - \tau_1^2$ and $\gamma = \tau_0^2 - \tau_1^2$. Eqn. 49 rewritten in terms of $\cos(\theta)$ using Chebyshev polynomials gives the attenuation ratio relative phase polynomial

$$\alpha T_{r_1}(\cos(\theta)) - \beta T_{r_0}(\cos(\theta)) + \gamma = 0, \tag{50}$$

where $T_n(x)$ is a Chebyshev polynomial of the first kind, defined by the trigonometric identity

$$T_n(x) = \cos(n\cos^{-1}(x)). \tag{51}$$

Chebyshev polynomials are equivalent to a cosine transform basis resampled using the function $\sqrt{1 - x^2}$ and are commonly used for polynomial fitting because they offer a good approximation to the min-max polynomial and avoid Runge's Phenomenon.[19] The first two Chebyshev polynomials are given by

$$T_1(x) = x, \tag{52}$$
$$T_2(x) = 2x^2 - 1. \tag{53}$$

By rewriting eqn. 49 using Chebyshev polynomials, a seemingly difficult equation has been reduced to a relatively simple polynomial. The order of the polynomial given by eqn. 50 is given by $\max(r_1, r_0)$. As the order increases, closed-form solutions become less useful and eventually impossible. The simplest possible case is for a frequency ratio of two to one. Given $r_1 = 2$ and $r_0 = 1$, eqn. 50 becomes

$$2\alpha\cos^2(\theta) - \beta\cos(\theta) - \alpha + \gamma = 0, \tag{54}$$

which has the solutions

$$\cos(\theta) \in \left\{1, -\frac{2\alpha - \beta}{2\alpha}\right\}. \tag{55}$$

4.3 Determining Relative Amplitude

The relative amplitude, b, is determined from eqn. 47 by substituting the second solution from eqn. 55, giving

$$b = \frac{\tau_0^2 - \cos(\theta) \pm \sqrt{(1 - \cos(\theta))(2\tau_0^2 - \cos(\theta) - 1)}}{1 - \tau_0^2} \tag{56}$$

$$\tag{57}$$

For valid (τ_0, τ_1), the positive square root variation gives $b \geq 1$ and the negative variation gives $b \in [0, 1]$ – the alternative values for b being reciprocals of each other. At most there are two unique candidate solution tuples (b, θ) and $(b, -\theta)$, since $(b^{-1}, -\theta)$ and (b^{-1}, θ) correspond to alternative normalisations of the same component returns.

4.4 Determining Absolute Phase and Amplitude

Given a known relationship between the two component returns, we now determine absolute phase and amplitude. This can be considered a denormalisation operation in the case where the unnormalised sum of the component returns and the normalised individual components are known. The component returns can be found for a candidate solution tuple (b, θ) by

$$
\begin{align}
\eta_0 &= \zeta \mathrm{e}^{2\pi jm/r_0} \tag{58}\\
\eta_1 &= b\eta_0 \mathrm{e}^{j\theta} \tag{59}
\end{align}
$$

where $\zeta \in \mathbb{C}$ is the undisambiguated* first component return calculated by

$$
\zeta = \frac{\xi_0^{\frac{1}{r_0}}}{|\xi_0|^{\frac{1}{r_0}-1}} \cdot \frac{|\Lambda_{\mathrm{f}}(b, r_0\theta)|^{\frac{1}{r_0}-1}}{\Lambda_{\mathrm{f}}(b, r_0\theta)^{\frac{1}{r_0}}} \tag{60}
$$

and the disambiguation constant, $m \in \mathbb{Z}$, is determined by writing out the second measurement (ξ_1) in terms of eqns. 58 and 59, giving

$$
\begin{align}
\xi_1 &= \frac{(\zeta \mathrm{e}^{2\pi jm/r_0})^{r_1}}{|\zeta \mathrm{e}^{2\pi jm/r_0}|^{r_1-1}} \cdot \Lambda(b, r_1\theta) \tag{61}\\
&= \frac{\zeta^{r_1}}{|\zeta|^{r_1-1}} \cdot \Lambda(b, r_1\theta) \mathrm{e}^{2\pi jmr_1/r_0} \tag{62}
\end{align}
$$

Valid candidate solution tuples can be discriminated from invalid tuples using m, which is an integer for a valid solution in the absence of noise.

When this algorithm is implemented on real data, a substantial reduction in error can be achieved by calculating (τ_0, τ_1) for each pixel and then constraining each value to a valid bound (e.g. $\tau_0 \in [0,1)$) before any further processing. This reduces the incidence of noise corrupted data values, that otherwise result in highly erroneous estimates of component returns.

5. PERFORMANCE ANALYSIS/DISCUSSION

In order to provide a brief analysis of the performance of the two new mixed pixel/multipath interference separation methods presented here, a series of simulatory and real-data experiments were performed. One of the major challenges in providing a quantitative metric for precision is the long-tailed nature of the error distributions. For example, fig. 2 gives the phase error distribution for a perfect single component return using the Cauchy method. While the noise distribution at first appears to be Gaussian, the data are extremely long-tailed and standard metrics like RMS error and variance have a tendency to produce non-representative results. There are several possible alternatives, including: plotting explicit histograms, albeit these can make direct comparisons difficult; plotting Cumulative Distribution Functions, such as utilised in fig. 3; and lastly, by fitting a parametric distribution more suited to representing the true shape of the error. We have chosen the latter approach; given that some of the error distributions appear nearly Cauchy distributed, a translated T-distribution has been fit using a Maximum Likelihood approach to the processed data. This distribution has three parameters: a mean, a width/spread coefficient and a degrees-of-freedom coefficient. The mean provides a measure of the accuracy of the estimates, thus any systematic perturbations introduced by the processing. The width coefficient is analogous to the standard deviation of a Gaussian distribution† and provides a measure of estimate precision, the error due to random noise. The remaining parameter, the degrees-of-freedom coefficient, encodes how long-tailed the distribution is.

---

*Before removal of the cyclic phase ambiguity in the component return (disambiguation).

†In the limit, as the number of degrees-of-freedom increases, width becomes equal to the standard deviation.
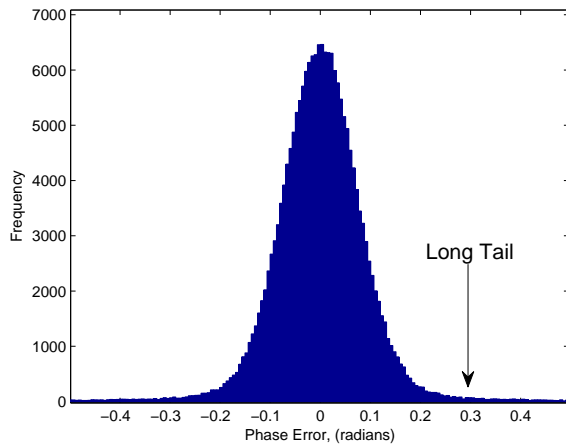
Figure 2: Error distribution of estimated phase in the case of a perfect single return at an SNR of 100:1 using the Cauchy method. Note that while the distribution appears similar to a Gaussian, it has extremely long, thin tails.
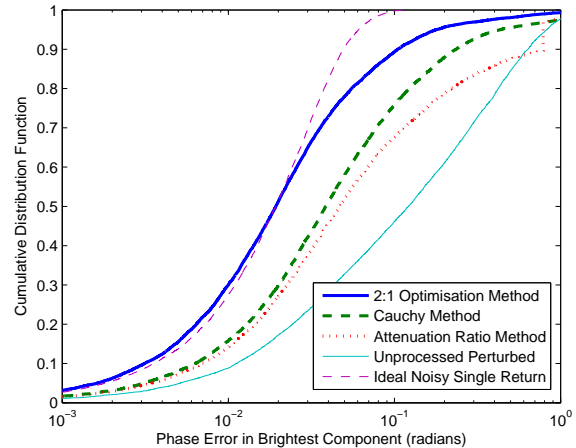

Figure 3: The Cumulative Distribution Function of phase error for different mixed pixel/multipath interference separation methods, assuming two component returns with distribution modelled by $\theta \sim \mathcal{U}(-\pi, \pi)$ and $b \sim \mathcal{U}(0, 1)$.

5.1 Simulatory Results

Assuming $\rho = 1$, a series of simulated measurements were produced at a specified SNR, assuming circularly symmetric complex Gaussian noise measured relative to the sum amplitude of the component returns. It was universally assumed that the brightest component return had a value of unity ($\eta_0 = 1$) so that the phase of each measurement naturally encoded the phase perturbation of the primary component return by the second component. All the component returns were simulated as perfect Dirac deltas, without distributed reflectivity. The measurements were processed by the attenuation ratio and Cauchy model algorithms presented in this paper. Fig. 2 gives an example of the type of long-tailed error distribution typically produced, while fig. 3 gives an overall comparison of the phase error performance of the two methods versus unprocessed noisy data and the 2:1 frequency ratio numerical optimisation based algorithm we presented in.[13] The latter graph clearly indicates that while both new methods substantially improve phase/range measurements in most cases, neither method produces results of the same quality as the 2:1 numerical method. In all cases, the processed data are particularly long-tailed compared to the reference case: the ideal noisy single component return.

A more detailed break-down of the behaviour of the two new algorithms is presented in fig. 4. Plotting systematic error as a function of relative phase shows that there is no improvement in systematic error until the component returns are separated by a certain phase, which is a function of SNR. Past this point, as the component returns begin to be separated out, precision is temporarily degraded, as indicated in fig. 4b. While the Cauchy algorithm only requires that the returns be separated by a particular relative phase, the attenuation ratio algorithm is also incapable of separating out component returns near $\theta = \pi$, resulting in similar breakdown behaviour. Fig. 4c shows the estimated attenuation coefficient, which is erroneous in the breakdown region, but otherwise accurate, and the amplitude of the brightest component return, which has nonzero phase angle in the breakdown region. This is contributed to by the inability to add additional non-linear constraints to the inversion, such as constraining coefficients to be real: this is one particular advantage to a numerical approach, allowing constraints which help improve resistance to noise. The final plot, fig. 4d, shows the estimated amplitude of the brightest component for different relative intensities, b. It appears that as the relative amplitude of the second component return influences the breakdown region: as the relative amplitude decreases, the threshold moves farther away from $\theta = 0$. Overall, the algorithms appear to trade off precision in order to achieve better accuracy; given that most commercial range-imaging cameras quote high precision, but are unable to produce equivalently high accuracy, this may be considered a suitable compromise in many circumstances.
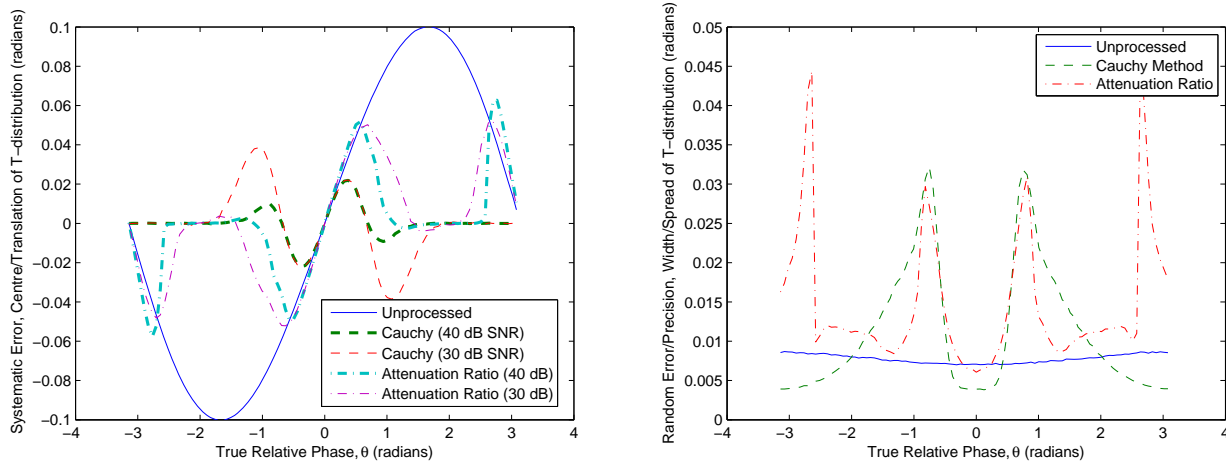
(a) Phase Accuracy

(b) Phase Precision
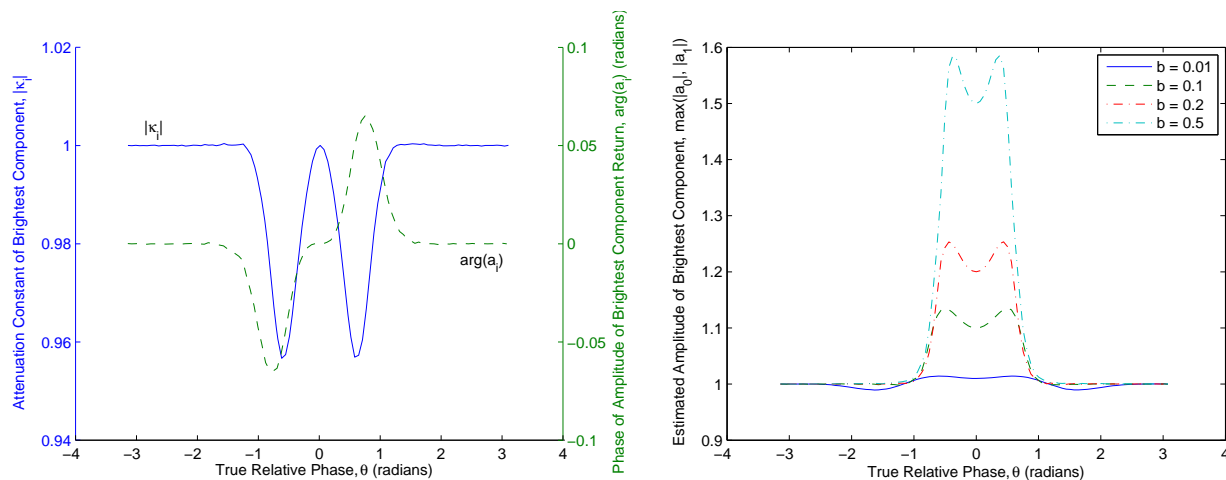
(c) Attenuation Coefficient and Phase of Amplitude
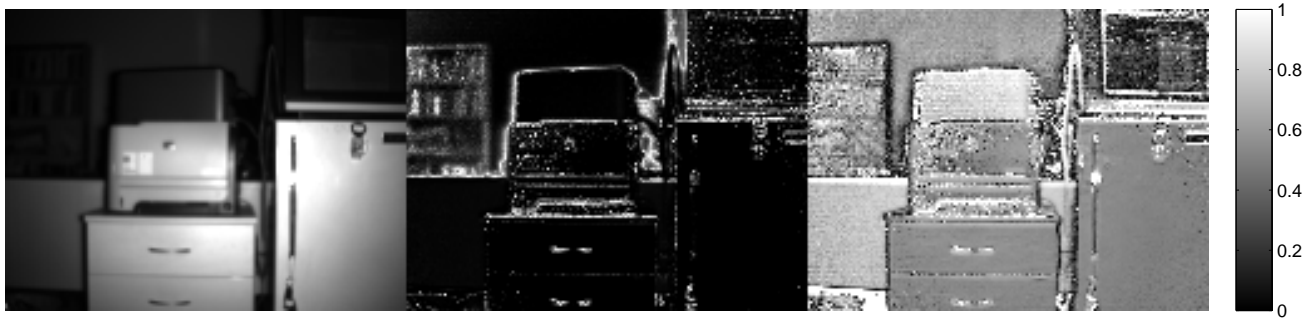

(d) Estimated Amplitude vs. Relative Amplitude

Figure 4: Estimates of component return properties as a function of relative phase. Unless otherwise specified, the results assume a SNR of 40 dB and $b = 0.1$ and plot the mean/translation parameter of the T-distribution model.

5.2 Real Data Results

A demonstration of the two algorithms was performed using the Canesta XZ-422 demonstrator camera. This was achieved by custom configuring five different modulation frequencies: 11 MHz, 22 MHz, 33 MHz and 44 MHz configurations were created using four equispaced phase steps. An additional modulation frequency was configured at 1.22 MHz in order to approximate a measurement of the zeroth spatial frequency. Ideally this measurement would have been made at an even lower modulation frequency to avert systematic errors due to partial cancellation, but it was not possible to achieve a more appropriate frequency, like 50 kHZ. One of the primary problems with the approach utilised for this paper is that of aliasing of correlation waveform harmonics, which are known to produce significant systematic errors.[14] One potential approach to partial amelioration is to utilise a larger number of phase steps, or apply harmonic cancellation.[20] While it is possible to calibrate for the errors introduced when there is only a single component return, for example by using a lookup table, it is not possible to calibrate when performing a mixed pixel/multipath separation operation without incorporating the calibration into the algorithm itself. Typical errors might be of the order of 30 milliradians phase error per measurement, which could easily result in a phase error several times that in the processed output of a separation

(a) Left to Right: Phase of Raw Data (11 MHz), Phase of Brightest Component, Phase of Second Component

(b) Left to Right: Amplitude of Raw Data, Estimated Relative Amplitude of Second Component b, Spread Coefficient

Figure 5: A scene processed by the Cauchy model algorithm. The spread coefficient ($|\kappa_i|$) has been rescaled so that black = 0.95 and white = 1.05.

algorithm[‡].

Images of a scene processed by the Cauchy algorithm are plotted in fig. 5. The algorithm has clearly detected regions subject to intra-camera scattering, primarily light from foreground objects scattered onto pixels measuring the background, however it has also introduced highly visible random noise. In particular, the bookcase in the background is very clearly displayed in the phase image, due to the varying reflectivity of books on the shelf. After processing, most of the perturbations have been removed. The phase of the darker components clearly indicates that it is light from the bright objects in front causing the perturbations. Unfortunately, the spread coefficient appears to be primarily overfitting noise, albeit one would expect there to be significant attenuation in regions around objects, such as at the edge of the bookcase, where the surface is nearly orthogonal to the camera and reflectivity would be spread over range. Fig. 6 is a plot of the boundary between two objects, showing how the Cauchy method has separated out the two objects within each pixel due to a combination of mixed pixels, defocus and intra-camera scattering; note that the darker return tends to be more noisy than the estimate of the brighter return. Fig. 7 shows the results of applying the attenuation ratio algorithm to the same scene. For the attenuation ratio algorithm, 22 MHz, 44 MHz and 1.22 MHz were utilised, so fig. 7a consists of a reference phase image at 22 MHz. Fig. 7 shows how α is a very good indicator of mixed returns, having a high correlation with estimates of b. Unfortunately, while some mixed pixels have been removed from the scene, the results are generally not as good as for the Cauchy algorithm.

6. CONCLUSION

In this paper we have presented two new closed-form methods for mixed pixel/multipath interference separation in AMCW lidar systems. One method models reflectivity as a Cauchy distribution over range and uses four measurements at different modulation frequencies to determine the amplitude, phase and reflectivity distribution of up to two component returns within each pixel. The other method uses attenuation ratios to determine the

---

[‡]There are also similar amplitude errors induced by aliasing.
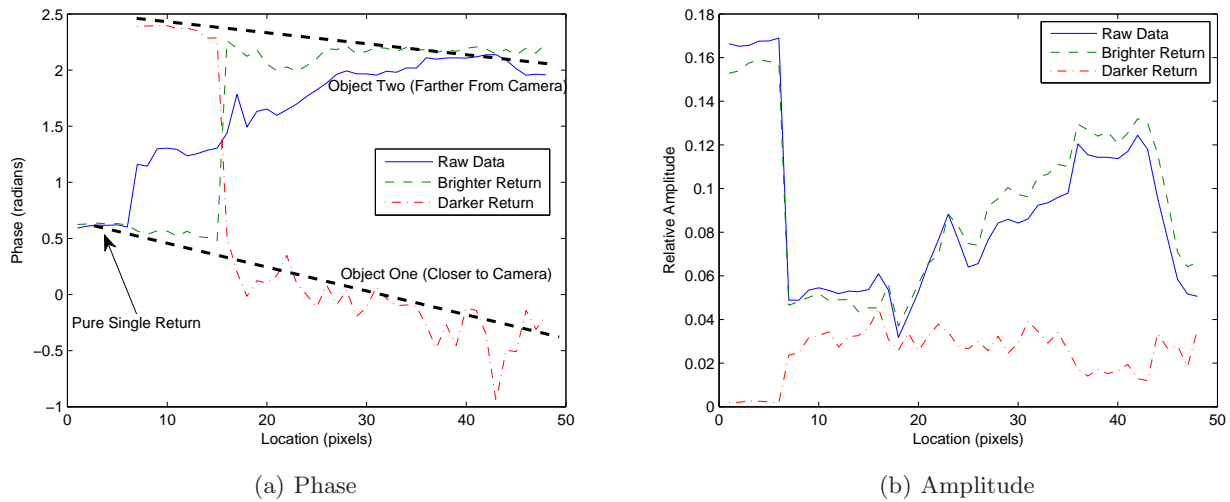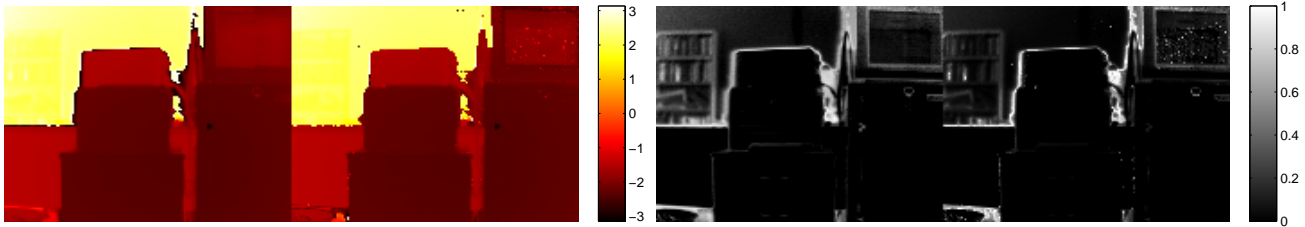
| (a) Phase | (b) Amplitude |

Figure 6: An example of the results from the Cauchy separation algorithm at an object boundary region subject to mild defocus and mixed pixels.

(a) Left to Right: Phase of Raw Data (22 MHz), Estimated Phase of Brightest Component

(b) Left to Right: Degree of Attenuation at r_0 (α), Estimated Relative Amplitude of Second Component (b)

Figure 7: The scene from fig. 5 processed by the attenuation ratio algorithm; note the close coupling between attenuation and estimated second component amplitude.

amplitude and phase of up to two component returns within each pixel. The methods were tested on both simulated and real data and shown to produce a significant improvement in overall error, although not quite as good as our previously published numerically based method. It was found that there is a particular relative phase threshold below which component returns cannot be separated; this threshold is a function of SNR and relative amplitude. Using simulation and real-data produced by the Canesta XZ-422 it was shown that both algorithms were capable of improving range images by trading worsened precision for improved accuracy.

References

[1] Kahlmann, T. and Ingensand, H., "Calibration and improvements of the high-resolution range-imaging camera Swissranger," in Videometrics VIII, **5665**, 144–155, SPIE, San Jose, CA, USA (2005).

[2] Godbaz, J. P., Dorrington, A. A., and Cree, M. J., "Understanding and ameliorating mixed pixels and multipath interference in AMCW lidar," in [Time-of-flight Range-Imaging Cameras], Remondino, F. and Stoppa, D., eds., Springer. In preparation.

[3] Karel, W., Ghuffar, S., and Pfeifer, N., "Quantifying the distortion of distance observations caused by scattering in time-of-flight range cameras," in International Archives of Photogrammetry, Remote Sensing and Spatial Information Sciences, Vol. XXXVIII, Part 5 Commision V Symposium, (2010).

[4] Falie, D. and Buzuloiu, V., "Noise characteristics of 3D time-of-flight cameras," in Proceedings of 2007 International Symposium on Signals, Circuits and Systems (ISSCS 2007), (2007).

[5] Guomundsson, S. A., Aanaes, H., and Larsen, R., "Environmental effects on measurement uncertainties of time-of-flight cameras," in *Proceedings of International Symposium on Signals, Circuits and Systems 2007, ISSCS 2007*, (2007).

[6] Hebert, M. and Krotkov, E., "3D measurements from imaging laser radars: How good are they?," IVC **10**, 170–178 (1992).

[7] Godbaz, J. P., Cree, M. J., and Dorrington, A. A., "Undue influence: Mitigating range-intensity coupling in AMCW 'flash' lidar using scene texture," *Proc. Image and Vision Computing New Zealand 2009*, 304–309 (2009).

[8] Falie, D., "Three-dimensional image corrections for time-of-flight cameras," IET Image Processing **5**(5), 523–529 (2011).

[9] Jamtsho, S. and Lichti, D., "Modeling scattering distortion of 3D range camera," in *Proceedings of the ISPRS Commission V Mid-Term Symposium on Close Range Image Measurement Techniques*, (2010).

[10] Fuchs, S., "Multipath interference compensation in time-of-flight camera images," in *Proceedings of the 2010 20th International Conference on Pattern Recognition*, ICPR '10, 3583–3586, IEEE Computer Society, Washington, DC, USA (2010).

[11] Godbaz, J. P., Cree, M. J., and Dorrington, A. A., "Mixed pixel return separation for a full-field ranger," in *Image and Vision Computing New Zealand 2008 (IVCNZ'08)*, 1–6 (2008).

[12] Godbaz, J. P., Cree, M. J., and Dorrington, A. A., "Multiple return separation for a full-field ranger via continuous waveform modelling," in *SPIE 7251 - Image Processing: Machine Vision Applications II*, **7251**, 72510T, SPIE, San Jose, CA, USA (2009).

[13] Dorrington, A. A., Godbaz, J. P., Cree, M. J., Payne, A. D., and Streeter, L. V., "Separating true range measurements from multi-path and scattering interference in commercial range cameras," in *Proc. of Three-Dimensional Imaging, Interaction, and Measurement*, 786404 (2011).

[14] Godbaz, J. P., Cree, M. J., and Dorrington, A. A., "Understanding and ameliorating non-linear phase and amplitude responses in AMCW lidar," Remote Sensing (2011). Pending publication.

[15] Dändliker, R., Salvadé, Y., and Zimmermann, E., "Distance measurement by multiple-wavelength interferometry Mesure de distance par interférométrie à plusieurs longueurs d'onde," Journal of Optics **29**, 105–114 (June 1998).

[16] Simpson, M. L., Cheng, M.-D., Dam, T. Q., Lenox, K. E., Price, J. R., Storey, J. M., Wachter, E. A., and Fisher, W. G., "Intensity-modulated, stepped frequency cw lidar for distributed aerosol and hard target measurements," Applied Optics **44**, 7210–7217 (2005).

[17] Hofton, M. A., Minster, J. B., and Blair, J. B., "Decomposition of laser altimeter waveforms," IEEE Transactions on Geoscience and Remote Sensing **38**, 1989–1996 (July 2000).

[18] Bishop, C. M., *Pattern Recognition and Machine Learning (Information Science and Statistics)*, Springer-Verlag New York, Inc., Secaucus, NJ, USA (2006).

[19] Press, W. H., Teukolsky, S. A., Vetterling, W. T., and Flannery, B. P., [Numerical Recipes 3rd Edition: The Art of Scientific Computing], Cambridge University Press, New York, NY, USA, 3 ed. (2007).

[20] Payne, A. D., Dorrington, A. A., Cree, M. J., and Carnegie, D. A., "Improved measurement linearity and precision for amcw time-of-flight range imaging cameras," Applied Optics **49**(23), 4392–4403 (2010).