

Swept-Angle Synthetic Wavelength Interferometry

Alankar Kotwal¹, Anat Levin², and Ioannis Gkioulekas¹

¹Carnegie Mellon University, ²Technion

Abstract

We present a new imaging technique, *swept-angle synthetic wavelength interferometry*, for full-field micron-scale 3D sensing. As in conventional synthetic wavelength interferometry, our technique uses light consisting of two narrowly-separated optical wavelengths, resulting in per-pixel interferometric measurements whose phase encodes scene depth. Our technique additionally uses a new type of light source that, by emulating spatially-incoherent illumination, makes interferometric measurements insensitive to aberrations and (sub)surface scattering, effects that corrupt phase measurements. The resulting technique combines the robustness to such corruptions of scanning interferometric setups, with the speed of full-field interferometric setups. Overall, our technique can recover full-frame depth at a lateral and axial resolution of $5\text{ }\mu\text{m}$, at frame rates of 5 Hz, even under strong ambient light. We build an experimental prototype, and use it to demonstrate these capabilities by scanning a variety of objects, including objects representative of applications in inspection and fabrication, and objects that contain challenging light scattering effects.

1. Introduction

Depth sensing is among the core problems of computer vision and computational imaging, with widespread applications in medicine, industry, and robotics. An array of techniques is available for acquiring depth maps of three-dimensional (3D) scenes at different scales. In particular, micrometer-resolution depth sensing, our focus in this paper, is important in biomedical imaging because biological features are often micron-scale, industrial fabrication and inspection of critical parts that must conform to their specifications (Figure 1), and robotics to handle fine objects.

Active illumination depth sensing techniques such as lidar, structured light, and correlation time-of-flight (ToF) cannot provide micrometer axial resolution. Instead, we focus on *interferometric* techniques that can achieve such resolutions. The operational principles and characteristics of interferometric techniques vary, depending on the type of active illumination and optical configurations they use.

The choice of illumination spectrum leads to techniques such as optical coherence tomography (OCT), which uses broadband illumination, and phase-shifting interferometry (PSI), which uses monochromatic illumination. We consider




Figure 1. **Applications of swept-angle SWI in industrial inspection and fabrication.** We show depth reconstructions for two scenes representative of these applications: millimeter-scale dents on an aircraft fuselage section, and a 3D-printed coin.

synthetic wavelength interferometry (SWI), which operates between these two extremes: By using illumination consisting of two narrowly-separated optical wavelengths, SWI provides a controllable trade-off between the large unambiguous depth range of OCT, and the large axial resolution of PSI. SWI can achieve micrometer resolution at depth ranges in the order of hundreds of micrometers.

The choice of optical configuration results in full-field versus scanning implementations, which offer different trade-offs. Full-field implementations acquire entire 2D depth maps, offering simultaneously fast operation and pixel-level lateral resolutions. However, full-field implementations are very sensitive to effects that corrupt depth estimation, such as imperfections in free-space optics (e.g., lens aberrations) and indirect illumination (e.g., subsurface scattering). By contrast, scanning implementations use beam steering to sequentially scan points in a scene and produce a 2D depth map. Scanning implementations offer robustness to depth corruption effects, through the use of fiber optics to reduce aberrations, and co-axial illumination and sensing to eliminate most indirect illumination. However, scanning makes acquiring depth maps with pixel-level lateral resolution and megapixel sizes impractically slow.

We develop a 3D sensing technique, *swept-angle synthetic wavelength interferometry*, that combines the complementary advantages of full-field and scanning implementations. We draw inspiration from previous work showing




Figure 2. **Reconstructing the eagle embossed on a \$20 bill.** The features on the eagle are raised $10\text{ }\mu\text{m}$ off the surface of the bill. The recovered depth shows fine details such as the gaps between the wings reconstructed with high lateral and axial resolution.

that the use of *spatially-incoherent illumination* in full-field interferometry mitigates aberration and indirect illumination effects. We design a light source that emulates spatial incoherence, by scanning *within exposure* a dichromatic point source at the focal plane of the illumination lens, effectively sweeping the illumination angle—hence *swept-angle*. We combine this light source with full-field SWI, resulting in a 3D sensing technique with the following performance characteristics: 5 Hz full-frame 2 Mpixel acquisition; $5\text{ }\mu\text{m}$ lateral and axial resolution; range and resolution tunability; and robustness to aberrations, ambient illumination, and indirect illumination. We build an experimental prototype, and use it to demonstrate these capabilities, as in Figure 2. We provide setup details, reconstruction code, and data in the supplement and project website.¹

Potential impact. Swept-angle SWI is relevant for critical applications, including industrial fabrication and inspection. In industrial fabrication, swept-angle SWI can be used to provide feedback during additive and subtractive manufacturing processes [67]. In industrial inspection, swept-angle SWI can be used to examine newly-fabricated or used in-place critical parts and ensure they comply with operational specifications. Swept-angle SWI, uniquely among 3D scanning techniques, offers a combination of features that are critical for both applications: First, high acquisition speed, which is necessary to avoid slowing down the manufacturing process, and to perform inspection efficiently. Second, micrometer lateral and axial resolution, which is necessary to detect critical defects. Third, robustness to aberrations and indirect illumination, which is necessary because of the materials used for fabrication, which often have strong subsurface scattering. Figure 1 showcases results representative of these applications: We scan a fuselage section from a Boeing aircraft, to detect critical defects such as scratches and bumps from collisions, at axial and lateral scales of a couple dozen micrometers. We also scan a coin pattern 3D-printed by a commercial material printer on a translucent material.

2. Related work

Depth sensing. There are many technologies for acquiring depth in computer vision. Passive techniques rely on scene appearance under ambient light, and use cues such as stereo [4, 30, 50], (de)focus [21, 31, 69], or shading [28, 34].

Table 1. Comparison of interferometric depth sensing techniques for millimeter-scale scenes. (‘f.f.’: full-field; ‘scan.’: scanning)

method	axial res	lateral res	depth range	acq time	robust
f.f. TD-OCT	✓	✓	✓	✗	✓
scan. FD-OCT	✗	✓	✓	✗	✓
scan. SS-OCT	✗	✓	✓	✗	✓
f.f. PSI	✓	✓	✗	✓	✗
scan. PSI	✗	✓	✗	✗	✓
f.f. SWI	✓	✓	✓	✓	✗
scan. SWI	✗	✓	✓	✗	✓
f.f. SA-SWI	✓	✓	✓	✓	✓

Active techniques inject controlled light in the scene, to overcome issues such as lack of texture and limited resolution. These include structured light [6, 24, 56, 65], impulse ToF [2, 19, 22, 23, 32, 39, 46, 52, 54, 62, 70, 71], and correlation ToF [3, 15, 16, 26, 27, 33, 42, 43, 58, 60, 66]. These techniques cannot easily achieve axial resolutions below hundreds of micrometers, placing them out of scope for applications requiring micrometer resolutions.

Optical interferometry. Interferometry is a classic wave-optics technique that measures the correlation, or *interference*, between two or more light beams that have traveled along different paths [29]. Most relevant to our work are three broad classes of interferometric techniques that provide depth sensing capabilities with different advantages and disadvantages, which we summarize in Table 1.

Phase-shifting interferometry (PSI) techniques [9, 37] use single-frequency illumination for nanometer-scale depth sensing. Similar to correlation ToF, they estimate depth by measuring the phase of a sinusoidal waveform. However, whereas correlation ToF uses external modulation to generate megahertz-frequency waveforms, PSI directly uses the terahertz-frequency light waveform. Unfortunately, the unambiguous depth range of PSI is limited to the illumination wavelength, typically around a micrometer.

Synthetic wavelength interferometry (SWI) (or *heterodyne interferometry*) techniques [7, 8, 10, 14, 44, 45] use illumination comprising multiple narrow spectral bands to *synthesize* waveforms at frequencies between the megahertz rates of correlation ToF and terahertz rates of PSI. These techniques allow control of the trade off between unambiguous depth range and resolution, by tuning the frequency of the *synthetic wavelength*. As our technique builds upon SWI, we detail its operational principles in Section 3.

¹https://imaging.cs.cmu.edu/swept_angle_swi

Lastly, *optical coherence tomography* (OCT) [20, 35, 41] uses broadband illumination to perform depth sensing analogously to impulse time-of-flight, by processing temporally-resolved (transient) responses. OCT decouples range and resolution, achieving unambiguous depth ranges up to centimeters, at micrometer axial resolutions. This flexibility comes at the cost of increased acquisition time, because of the need for axial (time-domain OCT) or lateral (Fourier-domain and swept-source OCT) scanning [13].

Mitigating indirect illumination effects. In their basic forms, active depth sensing techniques assume the presence of only direct illumination in the scene. *Indirect illumination* effects, such as interreflections and subsurface scattering, confound depth information. For example, in correlation ToF techniques (including PSI and SWI), indirect illumination results in incorrect phase, and thus depth, estimation, an effect also known as *multi-path interference*. Several techniques exist for *computationally* mitigating indirect illumination effects, using models of multi-bounce transport [18, 36, 49], sparse reconstruction [17, 38], multi-wavelength approaches [5], and neural approaches [48, 68].

Other techniques *optically* remove indirect illumination by *probing light transport* [57]. Such techniques include epipolar imaging [53, 56], high-spatial-frequency illumination [51, 61], and spatio-temporal coded illumination [1, 25, 55]. Similar probing capabilities are possible in interferometric systems, by exploiting the spatio-temporal coherence properties of the illumination [20, 41]. We adapt these techniques for robust micrometer-scale depth sensing.

3. Background on interferometry

The Michelson interferometer. Our optical setup is based on the classical Michelson interferometer (Figure 3(c)). The interferometer uses a beam splitter to divide collimated input illumination into two beams: one propagates toward the *scene arm*, and another propagates toward the *reference arm*—typically a planar mirror mounted on a translation stage that can vary the mirror’s distance from the beam splitter. After reflection, the two light beams recombine at the beam splitter and propagate toward the sensor.

We denote by l and $d(x)$ the distance from the beamsplitter of the reference mirror and the scene point that pixel x images, respectively. As l is a controllable parameter, we denote it explicitly. We denote by $\mathbf{u}_r(x, l)$ and $\mathbf{u}_s(x)$ the complex fields arriving at sensor pixel x from the reference and scene arms respectively. Then, the sensor measures,

$$I(x, l) = \underbrace{|\mathbf{u}_s(x)|^2 + |\mathbf{u}_r(x, l)|^2}_{\equiv \mathcal{I}(x, l)} + 2 \operatorname{Re}\{\underbrace{\mathbf{u}_s(x) \mathbf{u}_r^*(x, l)}_{\equiv \mathcal{C}(x, l)}\}. \quad (1)$$

The first two terms in Equation (1) are the intensities the sensor would measure if it were observing each of the two arms separately. We call their sum the *interference-free* image $\mathcal{I}$. The third term, which we call *interference*, is the real part of the complex *correlation* $\mathcal{C}$ between the reflected




Figure 3. **Synthetic wavelength interferometry.** (a-b) Collimated illumination from a point source emitting at two narrowly-separated wavelengths is injected into a Michelson interferometer. (c) As the reference mirror position l is scanned, each wavelength contributes to the interference a sinusoid with period equal to its wavelength. (d) The sum of these sinusoids has an envelope that is another sinusoid at a *synthetic wavelength* peaked at $l = d(x)$.

scene and reference fields. We elaborate on how to isolate the interference term from Equation (1) in Section 4.

Synthetic wavelength interferometry. SWI uses illumination comprising two distinct, but narrowly-separated, wavelengths that are incoherent with each other. We denote these wavelengths as λ and $\lambda/(1+\epsilon)$, corresponding to wavenumbers $\kappa \equiv 2\pi/\lambda$ and $(1+\epsilon)\kappa$, respectively. We assume that this illumination is injected into the interferometer as a collimated beam—for example, created by placing the outputs of two fiber-coupled single-frequency lasers at the focal plane of a lens, as in Figures 3 and 5. Then, we show in the supplement that the correlation $\mathcal{C}(x, l)$ equals

$$\mathcal{C}(x, l) = \exp(-2i\kappa(d(x) - l)) [1 + \exp(-2i\kappa\epsilon(d(x) - l))]. \quad (2)$$

The interference component of the camera intensity measurements in Equation (1) equals,

$$\operatorname{Re}(\mathcal{C}(x, l)) = 2 \sin(\kappa(2+\epsilon)(d(x) - l)) \sin(\kappa\epsilon(d(x) - l)) \quad (3)$$

$$\approx 2 \sin(2\kappa(d(x) - l)) \sin(\kappa\epsilon(d(x) - l)). \quad (4)$$

The approximation is accurate when $\epsilon \ll 1$, i.e., when the two wavelengths are close. We observe that, as a function of $d(x) - l$, the interference is the product of two sinusoids: first, a *carrier sinusoid* with *carrier wavelength* $\lambda_c \equiv \lambda/2$ and corresponding *carrier wavenumber* $\kappa_c \equiv 2\lambda$; second, an *envelope sinusoid* $\mathcal{E}$ with *synthetic wavelength* $\lambda_s \equiv \lambda/\epsilon$

and corresponding *synthetic wavenumber* $\kappa_s \equiv \epsilon\kappa$:

$$\mathcal{E}(x, l) \equiv \sin(\kappa_s(d(x) - l)). \quad (5)$$

Figure 3(d) visualizes $\text{Re}\{\mathcal{C}(x, l)\}$ and $\mathcal{E}(x, l)$. In practice, we measure only the squared amplitude of the envelope,

$$|\mathcal{E}(x, l)|^2 = \sin^2(\kappa_s(d(x) - l)) = \frac{1 - \cos(2\kappa_s(d(x) - l))}{2}. \quad (6)$$

From Equation (6), we see that SWI encodes scene depth $d(x)$ in the phase $\phi(d(x)) \equiv 2\kappa_s(d(x) - l)$ of the envelope sinusoid. We defer details on how to measure the envelope and estimate this phase until Section 4. We make two observations: First, SWI provides depth measurements at intervals of $[0, \lambda_s/2]$, and cannot disambiguate between depths differing by an integer multiple of $\lambda_s/2$. Second, the use of two wavelengths makes it possible to control the unambiguous depth range: by decreasing the separation ϵ between the two wavelengths, we increase the unambiguous depth range, at the cost of decreasing depth resolution.

Full-field and scanning interferometry. Figure 5 shows two types of Michelson interferometer setups that implement SWI: (a) a full-field interferometer; and (b) a scanning interferometer. We discuss their relative merits, which will motivate our proposed swept-angle interferometer setup.

Full-field interferometers (Figure 5(a)) use *free-space optics* to illuminate and image the entire field of view in the scene and reference arms. They also use a two-dimensional sensor to measure interference at all locations x . This enables fast interference measurements for all scene points at once, at lateral resolutions equal to the sensor pixel pitch.

Unfortunately, full-field interferometers are susceptible to phase corruption effects, as we visualize in Figure 4. Equation (5) assumes that the scene field is due to only the *direct* light path (solid line in Figure 4(a)), which produces a sinusoidal envelope with phase delay $d(x)$ (solid orange curve in Figure 4(b)). In practice, the scene field will include contributions from additional paths: First, *indirect* paths due to subsurface scattering (dashed lines in Figure 4(a)). Second, stray paths due to optical aberrations (Figure 4(c)). These paths have different lengths, and thus contribute to the envelope sinusoidal terms of different phases (dashed curves in Figure 4(b)). Their summation produces an overall sinusoidal envelope (black) with phase $d' \neq d(x)$, resulting in incorrect depth estimation.

Scanning interferometers (Figure 5(b)) use *fiber optics* (couplers, circulators, collimators) to generate a focused beam that illuminates only one point in the scene and reference arms. Additionally, they use a single-pixel sensor, focused at the same point. They also use steering optics (e.g., MEMS mirrors) to scan the focus point and capture interference measurements for the entire scene.

Scanning interferometers effectively mitigate the phase corruption effects in Figure 4. Because, at any given time, they only illuminate and image one point in the scene, they




Figure 4. Phase corruption effects. (a) A typical scene contains both direct (solid line) and indirect (dashed lines) light paths. (b) The direct path contributes a sinusoid with the correct phase (solid orange line). The indirect paths contribute sinusoids with incorrect phases (dashed lines). Their summation results in erroneous phase estimation (dark solid line). (c) Aberrations in free-space optics result in stray light paths that also contribute incorrect phases. (d) Images of the reference arm (mirror) visualize the effects of aberrations, and their mitigation using swept-angle illumination.

eliminate contributions from indirect paths such as those in Figure 4(a). Additionally, because the use of fiber optics minimizes aberrations, they eliminate stray paths such as those in Figure 4(c). Unfortunately, this robustness comes at the cost of having to use beam steering to scan the entire scene. This translates into long acquisition times, especially when it is necessary to measure depth at pixel-level lateral resolutions and at a sensor-equivalent field of view.

In the next section, we introduce an interferometer design that combines the fast acquisition and high lateral resolution of full-field interferometers, with the robustness to phase corruption effects of scanning interferometers.

4. Swept-angle illumination

In this section, we design a new light source for use with full-field SWI, to mitigate the phase corruption effects shown in Figure 4—indirect illumination, aberrations. Figure 5(c) shows our final *swept-angle* optical design.

Interferometry with spatially-incoherent illumination. Our starting point is previous findings on the use of *spatially-incoherent illumination* in full-field interferometric implementations. Specifically, Gkioulekas et al. [20] showed that using spatially incoherent illumination in a Michelson interferometer is equivalent to *direct-only probing*: this optically rejects indirect paths and makes the correlation term $\mathcal{C}$ depend predominantly on contributions from direct paths. Additionally, Xiao et al. [72] showed that using spatially-incoherent illumination makes the correlation term $\mathcal{C}$ insensitive to aberrations in the free-space optics.

Both Gkioulekas et al. [20] and Xiao et al. [72] realize spatially-incoherent illumination by replacing the point emitter in Figure 3 with an area emitter, e.g., LED or halogen lamp. Unfortunately, fundamental physics dictate that extending emission area of a light source is coupled with




Figure 5. **Synthetic wavelength interferometry setups.** (a) A full-field interferometer efficiently acquires full-frame depth, but is susceptible to phase corruptions due to aberrations and indirect illumination. (b) A scanning interferometer is robust to such corruptions, but requires slow lateral scanning. (c) A swept-angle full-field interferometer achieves both efficiency and robustness.




Figure 6. **The $\{4, 4\}$ -shift phase retrieval pipeline.** (a) We measure intensity at 16 reference positions, corresponding to 4 synthetic by 4 carrier subwavelength shifts. (b) and (c) For each synthetic sub-wavelength shift, we estimate *interference-free* and *envelope* images. (d) We denoise the envelope image using joint bilateral filtering. (e) We use 4-shift phase retrieval to estimate envelope phase and depth.

broadening its emission spectrum. This coupling was not an issue for prior work, which focused on OCT applications that require broadband illumination; but it makes spatially-incoherent light sources incompatible with SWI, which requires narrow-linewidth dichromatic illumination.

Emulating spatial coherence. To resolve this conundrum, we consider the following: Replacing the point emitter in Figure 3 with an area emitter changes the illumination, from a single collimated beam parallel to the optical axis, to the superposition of beams traveling along directions offset from the optical axis by angles $\theta \in [-\Theta/2, \Theta/2]$, where Θ depends on the emission area and lens focal length. We denote by $\mathbf{u}_s^\theta(x)$ and $\mathbf{u}_r^\theta(x, l)$ the complex fields resulting by each such beam reflecting on the scene and reference arms. Then, because different points on the area emitter are incoherent with each other, we can update Equation (1) as

$$I(x, l) = \int_{\theta} I^\theta(x, l) d\theta, \\ \mathcal{I}(x, l) = \int_{\theta} \mathcal{I}^\theta(x, l) d\theta, \quad \mathcal{C}(x, l) = \int_{\theta} \mathcal{C}^\theta(x, l) d\theta, \quad (7) \\ I^\theta(x, l) \equiv \underbrace{|\mathbf{u}_s^\theta(x)|^2 + |\mathbf{u}_r^\theta(x, l)|^2}_{\equiv \mathcal{I}^\theta(x, l)} + 2 \operatorname{Re} \{ \underbrace{\mathbf{u}_s^\theta(x) \mathbf{u}_r^{\theta*}(x, l)}_{\equiv \mathcal{C}^\theta(x, l)} \}.$$

We can interpret Equation (7) as follows: the image I measured using an area emitter equals the sum of the images that would be measured using independent point emitters spanning the emission area, each producing illumination at

an angle θ (and respectively for correlation $\mathcal{C}$).

With this interpretation at hand, we design the setup of Figure 5(c) to emulate a spatially-incoherent source suitable for SWI through time-division multiplexing. We use a galvo mirror to steer a narrow collimated beam of narrow-linewidth dichromatic illumination. We focus this beam through a relay lens at a point on the focal plane of the illumination lens. As we steer the beam direction, the focus point scans an area on the focal plane. In turn, each point location results in the injection of illumination at an angle θ in the interferometer. By scanning the focus point over a target effective emission area *within a single exposure*, we acquire image and correlation measurements that equal those of Equation (7). As this design sweeps the illumination angle, we term it *swept-angle illumination*, and its combination with full-field SWI *swept-angle synthetic wavelength interferometry*.

Comparison to Kotwal et al. [41]. The illumination module of the proposed setup in Figure 5(c) is an adaptation of the Fourier-domain redistributive projector of Kotwal et al. [41]. Compared to their setup, we do not need to use an amplitude electro-optic modulator, as we are only interested in emulating spatial incoherence (equivalent, direct-only probing). Additionally, their setup was designed for use with monochromatic illumination, to enable different light transport probing capabilities. By contrast, we use our setup with narrow-linewidth dichromatic illumination, to enable synthetic wavelength interferometry. More broadly,

our key insight is that swept-angle illumination can emulate extended-area emitters with arbitrary spectral profiles.

5. The $\{M, N\}$ -shift phase retrieval algorithm

We discuss our pipeline for acquiring and processing measurements to estimate the envelope phase, and thus depth. For this, in Figure 5(c), we use a nanometer-accuracy stage to vary the position l of the reference mirror.

Phase retrieval. We use the N -shift phase retrieval algorithm [9, 63, 64] to compute the envelope phase: We assume we have estimates of the envelope's squared amplitude $|\mathcal{E}(x, l_n)|^2$ at reference positions l_n corresponding to shifts by N -th fractions of the synthetic half-wavelength $\lambda_s/2$, $l_n = l + n\lambda_s/2N$, $n \in \{0, \dots, N-1\}$. Then, from Equation (6), we estimate the envelope phase as

$$\phi(d(x)) = \arctan \left[\frac{\sum_{n=0}^{N-1} |\mathcal{E}(x, l_n)|^2 \sin(2\pi n/N)}{\sum_{n=0}^{N-1} |\mathcal{E}(x, l_n)|^2 \cos(2\pi n/N)} \right], \quad (8)$$

and the depth (up to an integer multiple of $\lambda_s/2$) as

$$d(x) = l + \phi(d(x))/2\kappa_s. \quad (9)$$

Envelope estimation. Equation (8) requires estimates of $|\mathcal{E}(x, l_n)|^2$. For each l_n , we capture intensity measurements $I(x, l_n^m)$ at M reference positions corresponding to shifts by M -th fractions of the carrier wavelength, $l_n^m = l_n + m\lambda_c/M$, $m \in \{0, \dots, M-1\}$. As the carrier wavelength λ_c is half the optical wavelength λ and orders of magnitude smaller than the synthetic wavelength λ_s , the envelope and interference-free image remain approximately constant across shifts, $\mathcal{E}(x, l_n^m) \approx \mathcal{E}(x, l_n)$, $\mathcal{I}(x, l_n^m) \approx \mathcal{I}(x, l_n)$. Then, from Equations (1) and (4),

$$I(x, l_n^m) = \mathcal{I}(x, l_n) + 2 \sin(\kappa_c(d(x) - l_n^m)) \mathcal{E}(x, l_n). \quad (10)$$

We estimate the interference-free image and envelope as

$$\mathcal{I}(x, l_n) = 1/M \sum_{m=0}^{M-1} I(x, l_n^m), \quad (11)$$

$$|\mathcal{E}(x, l_n)|^2 = 1/2M \sum_{m=0}^{M-1} (I(x, l_n^m) - \mathcal{I}(x, l_n))^2. \quad (12)$$

Equations (11), (12), (8), and (9) complete our phase estimation pipeline, which we summarize in Figure 6. We note that the runtime of this pipeline is negligible relative to acquisition time. We provide algorithmic details and pseudocode in the supplement. We call this pipeline the $\{M, N\}$ -shift phase retrieval algorithm. The parameters M and N are design parameters that we can fine-tune to control a trade-off between acquisition time and depth accuracy: As we increase M and N , the number $M \cdot N$ of images we capture also increases, and depth estimates become more robust to noise. The theoretical minimum number of images is achieved using $\{3, 3\}$ shifts (9 images). In practice, we found $\{4, 4\}$ shifts to be robust. We evaluate different $\{M, N\}$ combinations in the supplement.

Dealing with speckle. Interference in non-specular scenes takes the form of *speckle*, a high-frequency pseudo-random pattern (Figure 6(a)). This can result in noisy envelope,




Figure 7. **Effect of swept-angle scanning.** Swept-angle illumination helps mitigate speckle noise and subsurface scattering effects.

and thus phase and depth, estimates. As Figures 4(d) and 7 show, swept-angle illumination greatly mitigates speckle effects. We can further reduce the impact of speckle by denoising estimated quantities with a low-pass filter (e.g., Gaussian) [20, 41]. Alternatively, to avoid blurring image details, we can use joint bilateral filtering [12, 59] with a guide image of the scene under ambient light. Empirically, we found it better to blur the envelope estimates of Equation (12) before using them in Equation (8) (Figure 6(d)), compared to blurring the phase or depth estimates.

6. Experiments

Experimental prototype. For all our experiments, we use the experimental prototype in Figure 5(c). Our prototype comprises: two distributed-Bragg-reflector lasers (wavelengths 780 nm and 781 nm, power 45 mW, linewidth 1 MHz); a galvo mirror pair; a translation stage (resolution 10 nm); two compound macro lenses (focal length 200 mm); a CCD sensor (pixel pitch 3.7 μm , 3400×2700 pixels); and other free-space optics and mounts. We provide a detailed component list and specifications, calibration instructions, and other information in the supplement.

We use the following experimental specifications: The reproduction ratio is 1:1, the field of view is 12.5 mm $\times$ 10 mm, and the working distance is 400 mm. The unambiguous depth range is approximately 500 μm . We use $\{4, 4\}$ -shifts (16 images), and a minimum per-image exposure time of 10 ms, resulting at a frame rate of 5 Hz.

Depth recovery on challenging scenes. Figures 8 and 2 show scans of many challenging scenes. We scan materials ranging from rough metallic (coins), to diffuse (music box, pill), to highly-scattering (soap and chocolate). Most scenes include fine features requiring high lateral and axial resolution (music box, business card, US quarter, dollar bill). We compare scan results using swept-angle SWI (with bilateral or Gaussian filtering), versus conventional full-field SWI (Figure 5(a), emulated by deactivating the angle-sweep galvo). In all scenes, using swept-angle illumination greatly improves reconstruction quality. The difference is more pronounced in scenes with strong subsurface scattering (music box, chocolate, soap). Even in metallic scenes where there is no subsurface scattering (coins), the use of swept-angle illumination still improves reconstruction quality, by helping mitigate aberration artifacts. Bilateral filtering further improves lateral detail compared to Gaussian filtering.

Comparison with scanning SWI. A direct comparison of




Figure 8. **Depth reconstruction.** Depth maps (left) and surface renderings (right) acquired using full-field SWI with: (c) swept-angle scanning and bilateral filtering; (d) swept-angle scanning and Gaussian filtering; (e) no swept-angle scanning and with Gaussian filtering.

swept-angle SWI with scanning SWI is challenging, because of the differences between scanning and full-field setups (Figure 5, (b) versus (a), (c)). To qualitatively assess their relative performance, in Figure 9 we use the following experimental protocol: We downsample the depth map from our technique, by a factor of 35 in each dimension. This is approximately equal to the number of depths points a scanning SWI system *operating in point-to-point mode* would acquire at the same total exposure time as ours, when equipped with beam-steering optics that operate at a 30 kHz scan rate—we detail the exact calculation in the supplement. We then use joint bilateral upsampling [40] with the same guide image as in our technique, to upsample the downsampled depth map back to its original resolution. We observe that, due to the sparse set of points scanning SWI can acquire, the reconstructions miss fine features such as the hair on the quarter and letters on the business card.

Scanning SWI systems can also operate in resonant mode, typically using Lissajous scanning to raster scan the field of

view [47]. Resonant mode enables much faster operation than the point-to-point mode we compare against. However, Lissajous scanning results in non-rectilinear sampling of the image plane, as scan lines smear across multiple image rows (when the fast axis is horizontal). At high magnifications, this produces strong spatial distortions and artifacts, making resonant mode operation unsuitable for applications that require micrometer lateral resolution. Our technique also uses resonant-mode Lissajous scanning, but to raster scan the focal plane of the illumination lens (Figure 5(d)), and not the image plane. Instead, our technique uses a two-dimensional sensor to sample the image plane, thus completely avoiding the artifacts above and achieving micrometer lateral resolution. The supplement discusses other challenges in achieving high lateral resolution with scanning SWI.

Additionally, our technique is orders of magnitude faster than resonant-mode scanning SWI at the same lateral resolution. This is because, whereas scanning SWI must scan the image plane at the target (pixel-level) lateral resolution, our




Figure 9. **Comparison with upsampled scanning SWI.** We emulate scanning SWI by downsampling our swept-angle SWI depth. We then bilaterally upsample it to the original resolution.

Table 2. **Depth accuracy.** MedAE is the median absolute error between ground truth and estimated depth. Kernel width is the lateral size of the speckle blur filter. All quantities are in μm .

kernel width	with swept-angle		w/o swept-angle	
	RMSE	MedAE	RMSE	MedAE
7	8.2	4.8	18.9	13.2
15	5.1	3.6	11.2	9.5
21	2.0	1.6	10.5	7.3
30	1.6	1.0	11.1	6.7

technique remains effective even when scanning the focal plane at much lower resolutions. Figure 10 shows images of the focal plane during scanning, taken with the measurement camera. The dark regions (insets) show that scanning resolution is much lower than pixel-level resolution.

Axial resolution. To quantify the axial resolution of our technique, we perform the following experiment. We use a second nanometer-accurate translation stage to place the chocolate scene from Figure 8 at different depths from the camera, at increments of $1\text{ }\mu\text{m}$. We choose this scene because it has strong sub-surface scattering. We then compare how well full-field SWI with and without swept-angle illumination can track the scene depth. We perform this experiment using Gaussian filtering with different kernel sizes, to additionally quantify lateral resolution. The results in Table 2 show that we can achieve an axial resolution of approximately $5\text{ }\mu\text{m}$ and $1\text{ }\mu\text{m}$, for kernel sizes $7\text{ }\mu\text{m}$ and $30\text{ }\mu\text{m}$, respectively. We can trade off lateral for axial resolution, by increasing kernel size. Figure 8 shows that this trade-off can become more favorable with bilateral filtering.

Additional experiments. In the supplement: 1) *Comparisons with full-field OCT.* We show that swept-angle SWI can approximate the reconstruction quality of OCT, despite being $50\times$ faster. 2) *Tunable depth range.* We show that, by adjusting the wavelength separation between the two lasers, swept-angle SWI can scan scenes with depth range 16 mm at resolution $50\text{ }\mu\text{m}$, at significantly better quality than con-



Figure 10. **Lissajous scanning.** Image and insets of the illumination lens' focal plane, showing the Lissajous scanning pattern.

ventional full-field SWI. 3) *Robustness to ambient lighting.* We show that the use of near-monochromatic illumination makes swept-angle SWI robust to strong ambient light, even at 10% signal-to-background ratios. 4) *Effect of scanning resolution and $\{M, N\}$ settings.* We show that, by adjusting the focal plane scanning resolution and $\{M, N\}$ parameters for phase retrieval, swept-angle SWI can trade off between acquisition time versus reconstruction quality.

7. Limitations and conclusion

Phase wrapping. SWI determines depth only up to an integer multiple of the synthetic half-wavelength $\lambda_s/2$. Equivalently, all depth values are *wrapped* to $[0, \lambda_s/2]$, even if the depth range is greater. To mitigate phase wrapping and extend unambiguous depth range, it is common to capture measurements at multiple synthetic wavelengths, and use them to *unwrap* the phase estimate [3, 7, 8, 11]. Theoretically two synthetic wavelengths are enough to uniquely determine depth, but in practice phase unwrapping techniques are very sensitive to noise [25]. Combining phase unwrapping with our technique is an interesting future direction.

Toward real-time operation. Our prototype acquires measurements at a frame rate of 5 Hz, due to the need for 10 ms per-frame exposure time. If we use a stronger laser to reduce this time, the main speed bottleneck will be the need to perform phase shifts by physically translating the reference mirror. We can mitigate this bottleneck using faster translation stages (e.g., fast microscopy stages).

Conclusion. We presented a technique for fast depth sensing at micron-scale lateral and axial resolutions. Our technique, swept-angle synthetic wavelength interferometry, combines the complementary advantages of full-field interferometry (speed, pixel-level lateral resolution) and scanning interferometry (robustness to aberrations and indirect illumination). We demonstrated these advantages by scanning multiple scenes with fine geometry and strong subsurface scattering. We expect our results to motivate applications of swept-angle SWI in areas such as biomedical imaging, and industrial fabrication and inspection (Figure 1).

Acknowledgments. We thank Sudershan Boovaraghavan, Yuvraj Agrawal, Arpit Agarwal, Wenzhen Yuan from CMU, and Veniamin V. Stryzheus, Brian T. Miller from The Boeing Company, who provided the samples for the experiments in Figure 1. This work was supported by NSF awards 1730147, 2047341, 2008123 (NSF-BSF 2019758), and a Sloan Research Fellowship for Ioannis Gkioulekas.

References

- [1] Supreeth Achar, Joseph R Bartels, William L Whittaker, Kiriakos N Kutulakos, and Srinivasa G Narasimhan. Epipolar time-of-flight imaging. *ACM TOG*, 2017. 3
- [2] Brian Aull. 3d imaging with geiger-mode avalanche photodiodes. *Opt. Photon. News*, 2005. 2
- [3] Seung-Hwan Baek, Noah Walsh, Ilya Chugunov, Zheng Shi, and Felix Heide. Centimeter-wave free-space neural time-of-flight imaging. *ACM Transactions on Graphics (TOG)*, 2022. 2, 8
- [4] Stephen T. Barnard and William B. Thompson. Disparity analysis of images. *IEEE TPAMI*, 1980. 2
- [5] Ayush Bhandari, Achuta Kadambi, Refael Whyte, Christopher Barsi, Micha Feigin, Adrian Dorrington, and Ramesh Raskar. Resolving multipath interference in time-of-flight imaging via modulation frequency diversity and sparse regularization. *Optics Letters*, 2014. 3
- [6] Tongbo Chen, Hans-Peter Seidel, and Hendrik P. A. Lensch. Modulated phase-shifting for 3d scanning. *IEEE/CVF CVPR*, 2008. 2
- [7] Yeou-Yen Cheng and James C. Wyant. Two-wavelength phase shifting interferometry. *Applied Optics*, 1984. 2, 8
- [8] Yeou-Yen Cheng and James C. Wyant. Multiple-wavelength phase-shifting interferometry. *Applied Optics*, 1985. 2, 8
- [9] Peter de Groot. *Phase Shifting Interferometry*. 2011. 2, 6
- [10] Peter de Groot and John McGarvey. Chirped synthetic-wavelength interferometry. *Optics Letters*, 1992. 2
- [11] David Droschel, Dirk Holz, and Sven Behnke. Multi-frequency phase unwrapping for time-of-flight cameras. *IEEE/RSJ IROS*, 2010. 8
- [12] Elmar Eisemann and Frédo Durand. Flash photography enhancement via intrinsic relighting. *ACM transactions on graphics (TOG)*, 2004. 6
- [13] Adolf F Fercher, Wolfgang Drexler, Christoph K Hitzenberger, and Theo Lasser. Optical coherence tomography-principles and applications. *Reports on progress in physics*, 2003. 3
- [14] Adolf F. Fercher, H Z. Hu, and U. Vry. Rough surface interferometry with a two-wavelength heterodyne speckle interferometer. *Applied Optics*, 1985. 2
- [15] Richard Ferriere, Johann Cussey, and John M. Dudley. Time-of-flight range detection using low-frequency intensity modulation of a cw laser diode: application to fiber length measurement. *Optical Engineering*, 2008. 2
- [16] Angel Flores, Craig Robin, Ann Lanari, and Iyad Dajani. Pseudo-random binary sequence phase modulation for narrow linewidth, kilowatt, monolithic fiber amplifiers. *Optics Express*, 2014. 2
- [17] Daniel Freedman, Eyal Krupka, Yoni Smolin, Ido Leichter, and Mirko Schmidt. Sra: Fast removal of general multipath for tof sensors, 2014. 3
- [18] Stefan Fuchs. Multipath interference compensation in time-of-flight camera images. *ICPR*, 2010. 3
- [19] Genevieve Gariepy, Nikola Krstajić, Robert Henderson, Chunyong Li, Robert R Thomson, Gerald S Buller, Barmak Heshmat, Ramesh Raskar, Jonathan Leach, and Daniele Faccio. Single-photon sensitive light-in-flight imaging. *Nature Comm.*, 2015. 2
- [20] Ioannis Gkioulekas, Anat Levin, Frédo Durand, and Todd Zickler. Micron-scale light transport decomposition using interferometry. *ACM TOG*, 2015. 3, 4, 6
- [21] Pavel Grossmann. Depth from focus. *PRL*, 1987. 2
- [22] Anant Gupta, Atul Ingle, and Mohit Gupta. Asynchronous single-photon 3d imaging. In *ICCV*, 2019. 2
- [23] Anant Gupta, Atul Ingle, Andreas Velten, and Mohit Gupta. Photon-flooded single-photon 3d cameras. *IEEE/CVF CVPR*, 2019. 2
- [24] Mohit Gupta, Amit Agrawal, Ashok Veeraraghavan, and Srinivasa G Narasimhan. Structured light 3d scanning in the presence of global illumination. *IEEE/CVF CVPR*, 2011. 2
- [25] Mohit Gupta, Shree K Nayar, Matthias B Hullin, and Jaime Martin. Phasor imaging: A generalization of correlation-based time-of-flight imaging. *ACM TOG*, 2015. 3, 8
- [26] Mohit Gupta, Andreas Velten, Shree K. Nayar, and Eric Breibach. What are optimal coding functions for time-of-flight imaging? *ACM TOG*, 2018. 2
- [27] Felipe Gutierrez-Barragan, Syed Azer Reza, Andreas Velten, and Mohit Gupta. Practical coding function design for time-of-flight imaging. *IEEE/CVF CVPR*, 2019. 2
- [28] Yudeog Han, Joon-Young Lee, and In So Kweon. High quality shape from a single rgb-d image under uncalibrated natural illumination. *IEEE ICCV*, 2013. 2
- [29] Parameswaran Hariharan. *Optical interferometry*. Elsevier, 2003. 2
- [30] Richard Hartley and Andrew Zisserman. *Multiple View Geometry in Computer Vision*. 2004. 2
- [31] Caner Hazirbas, Sebastian Georg Soyer, Maximilian Christian Staab, Laura Leal-Taixé, and Daniel Cremers. Deep depth from focus, 2018. 2
- [32] Felix Heide, Steven Diamond, David B. Lindell, and Gordon Wetzstein. Sub-picosecond photon-efficient 3d imaging using single-photon sensors. *Scientific Reports*, 2018. 2
- [33] Felix Heide, Matthias B Hullin, James Gregson, and Wolfgang Heidrich. Low-budget transient imaging using photonic mixer devices. *ACM TOG*, 2013. 2
- [34] Berthold Klaus Paul Horn. Shape from shading: A method for obtaining the shape of a smooth opaque object from one view. Technical report, 1970. 2
- [35] David Huang, Eric A Swanson, Charles P Lin, Joel S Schuman, William G Stinson, Warren Chang, Michael R Hee, Thomas Flotte, Kenton Gregory, Carmen A Puliafito, and James G Fujimoto. Optical coherence tomography. *Science*, 1991. 3
- [36] David Jimenez, Daniel Pizarro, Manuel Mazoa, and Sira Palazuelos. Modeling and correction of multipath interference in time of flight cameras. *Image and Vision Computing*, 2014. 3
- [37] Jon L. Johnson, Timothy D. Dorney, and Daniel M. Mittleman. Enhanced depth resolution in terahertz imaging using phase-shift interferometry. *Applied Physics Letters*, 2001. 2
- [38] Achuta Kadambi, Refael Whyte, Ayush Bhandari, Lee Streeter, Christopher Barsi, Adrian Dorrington, and Ramesh

- Raskar. Coded Time of Flight Cameras: Sparse Deconvolution to Address Multipath Interference and Recover Time Profiles. *ACM TOG*, 2013. 3
- [39] Ahmed Kirmani, Dheera Venkatraman, Dongeek Shin, Andrea Colaço, Franco N. C. Wong, Jeffrey H. Shapiro, and Vivek K Goyal. First-photon imaging. *Science*, 2014. 2
- [40] Johannes Kopf, Michael F. Cohen, Dani Lischinski, and Matt Uyttendaele. Joint bilateral upsampling. *ACM Trans. Graph.*, 26(3):96–es, jul 2007. 7
- [41] Alankar Kotwal, Anat Levin, and Ioannis Gkioulekas. Interferometric transmission probing with coded mutual intensity. *ACM TOG*, 2020. 3, 5, 6
- [42] Robert Lange and Peter Seitz. Solid-state time-of-flight range camera. *IEEE JQE*, 2001. 2
- [43] Robert Lange, Peter Seitz, Alice Biber, and Stefan Lauxtermann. Demodulation pixels in ccd and cmos technologies. *SPIE*, 2000. 2
- [44] Fengqiang Li, Florian Willomitzer, Prasanna Rangarajan, Mohit Gupta, Andreas Velten, and Oliver Cossairt. Sh-tof: Micro resolution time-of-flight imaging with superheterodyne interferometry. *IEEE ICCP*, 2018. 2
- [45] Fengqiang Li, Joshua Yablon, Andreas Velten, Mohit Gupta, and Oliver Cossairt. High-depth-resolution range imaging with multiple-wavelength superheterodyne interferometry using 1550-nm lasers. *Applied Optics*, 2017. 2
- [46] David B Lindell, Matthew O’Toole, and Gordon Wetzstein. Single-photon 3d imaging with deep sensor fusion. *ACM TOG*, 2018. 2
- [47] Xiaomeng Liu, Kristofer Henderson, Joshua Rego, Suren Jayasuriya, and Sanjeev Koppal. Dense lissajous sampling and interpolation for dynamic light-transport. *Optics Express*, 2021. 7
- [48] Julio Marco, Quercus Hernandez, Adolfo Muñoz, Yue Dong, Adrian Jarabo, Min H. Kim, Xin Tong, and Diego Gutierrez. Deeptof: Off-the-shelf real-time correction of multipath interference in time-of-flight imaging. *ACM TOG*, 2017. 3
- [49] Nikhil Naik, Achuta Kadambi, Christoph Rhemann, Shahram Izadi, Ramesh Raskar, and Sing Bing Kang. A light transport model for mitigating multipath interference in time-of-flight sensors. *IEEE/CVF CVPR*, 2015. 3
- [50] Lazaros Nalpantidis, Georgios Sirakoulis, and Antonios Gasteratos. Review of stereo vision algorithms: From software to hardware. *Int. J. Optomechatronics*, 2008. 2
- [51] Shree Nayar, Gurunandan Krishnan, Michael Grossberg, and Ramesh Raskar. Fast separation of direct and global components of a scene using high frequency illumination. *ACM TOG*, 2006. 3
- [52] Cristiano Niclass, Alexis Rochas, Pierre. Besse, and Edoardo Charbon. Design and characterization of a cmos 3-d image sensor based on single photon avalanche diodes. *IEEE JSSC*, 2005. 2
- [53] Matthew O’Toole, Supreeth Achar, Srinivasa G Narasimhan, and Kiriakos N Kutulakos. Homogeneous codes for energy-efficient illumination and imaging. *ACM TOG*, 2015. 3
- [54] Matthew O’Toole, Felix Heide, David B Lindell, Kai Zang, Steven Diamond, and Gordon Wetzstein. Reconstructing transient images from single-photon sensors. *IEEE/CVF CVPR*, 2017. 2
- [55] Matthew O’Toole, Felix Heide, Lei Xiao, Matthias B Hullin, Wolfgang Heidrich, and Kiriakos N Kutulakos. Temporal Frequency Probing for 5D Transient Analysis of Global Light Transport. *ACM TOG*, 2014. 3
- [56] Matthew O’Toole, John Mather, and Kiriakos N. Kutulakos. 3d shape and indirect appearance by structured light transport. *IEEE TPAMI*, 2016. 2, 3
- [57] Matthew O’Toole, Ramesh Raskar, and Kiriakos N Kutulakos. Primal-dual Coding to Probe Light Transport. *ACM TOG*, 2012. 3
- [58] Andrew Payne, A Dorrington, and Michael Cree. Illumination waveform optimization for time-of-flight range imaging cameras. *SPIE*, 2011. 2
- [59] Georg Petschnigg, Richard Szeliski, Maneesh Agrawala, Michael Cohen, Hugues Hoppe, and Kentaro Toyama. Digital photography with flash and no-flash image pairs. *ACM transactions on graphics (TOG)*, 2004. 6
- [60] Dario Piatti, Fabio Remondino, and David Stoppa. *State-of-the-Art of TOF Range-Imaging Sensors*. 2013. 2
- [61] Dikpal Reddy, Ravi Ramamoorthi, and Brian Curless. Frequency-space Decomposition and Acquisition of Light Transport Under Spatially Varying Illumination. *ECCV*, 2012. 3
- [62] Alexis Rochas, Michael Gosch, Alexandre Serov, Pierre Besse, Rade S. Popovic, T. Lasser, and Rudolph Rigler. First fully integrated 2-d array of single-photon detectors in standard cmos technology. *IEEE PTL*, 2003. 2
- [63] Eric W Rogala and Harrison H Barrett. Phase-shifting interferometry and maximum-likelihood estimation theory. *Applied optics*, 36(34):8871–8876, 1997. 6
- [64] Eric W Rogala and Harrison H Barrett. Phase-shifting interferometry and maximum-likelihood estimation theory. ii. a generalized solution. *Applied optics*, 37(31):7253–7258, 1998. 6
- [65] Daniel Scharstein and Richard Szeliski. High-accuracy stereo depth maps using structured light. *IEEE/CVF CVPR*, 2003. 2
- [66] Rudolf Schwarte, Zhanping Xu, Horst-Guenther Heinol, Joachim Olk, Ruediger Klein, Bernd Buxbaum, Helmut Fischer, and Juergen Schulte. New electro-optical mixing and correlating sensor: facilities and applications of the photonic mixer device (PMD). In *Sensors, Sensor Systems, and Sensor Data Processing*, 1997. 2
- [67] Pitchaya Sitthi-Amorn, Javier E Ramos, Yuwang Wang, Joyce Kwan, Justin Lan, Wenshou Wang, and Wojciech Matusik. Multifab: a machine vision assisted platform for multi-material 3d printing. *Acm Transactions on Graphics (Tog)*, 2015. 2
- [68] Shuochen Su, Felix Heide, Gordon Wetzstein, and Wolfgang Heidrich. Deep end-to-end time-of-flight imaging. *CVPR*, 2018. 3
- [69] Murali Subbarao and Gopal Surya. Depth from defocus: A spatial domain approach. *IJCV*, 1994. 2
- [70] Andreas Velten, Thomas Willwacher, Otkrist Gupta, Ashok Veeraraghavan, Mounsi G. Bawendi, and Ramesh Raskar. Recovering three-dimensional shape around a corner using ultrafast time-of-flight imaging. *Nature Comm.*, 2012. 2

- [71] Federica Villa, Rudi Lussana, Danilo Bronzi, Simone Tisa, Alberto Tosi, Franco Zappa, Alberto Dalla Mora, Davide Contini, Daniel Durini, Sasha Weyers, and Werner Brockherde. Cmos imager with 1024 spads and tdcs for single-photon timing and 3-d time-of-flight. *IEEE JSTQE*, 2014. 2
- [72] Peng Xiao, Mathias Fink, and A Claude Boccara. Full-field spatially incoherent illumination interferometry: a spatial resolution almost insensitive to aberrations. *Optics letters*, 2016. 4