SDR-ToF: Million Optical Depth Samples per Second Using Software-Defined Radio

ACM Reference Format:
. 2026. SDR-ToF: Million Optical Depth Samples per Second Using Software-Defined Radio. *ACM Trans. Graph.* 44, 400, Article 400 (December 2026), 6 pages. https://doi.org/10.1111/11111111

Contents

1 SUPPLEMENTARY PACKAGE CONTENTS

The supplementary results are best visualized through the included local web viewer. To use it, unzip the supplementary folder, start a local server from the unzipped directory, and open the corresponding localhost URL in a browser:

```
python -m http.server 2127 -b 0.0.0.0
```

For the default port, open:

```
http://localhost:2127
```

If a different port is printed by the server, replace `2127` with that port number.

The supplementary package contains the following materials:

- `audio/`: audio files for the visual microphone results.
- `dynamic_depth_results/`: dynamic depth results, including 1D and 2D scans.
- `figures/`: figures and velocimetry results.
- `static_depth_ply/`: static depth point clouds in PLY format.
- `method_video_sdr_tof.mp4`: method video for the SDR-ToF system.
- Anonymous code repository: https://anonymous.4open.science/r/sdr-tof-2127s1/.

Author's Contact Information:

ACM 1557-7368/2026/12-ART400
https://doi.org/10.1111/11111111

2 ABLATION: INTENSITY VERSUS PHASE READOUT

A vibrating surface changes both the optical path length and the surface normal. SDR-ToF measures the former; the latter deflects the return across the detector aperture and reorganizes the speckle pattern, modulating the received intensity. A natural concern is whether intensity fluctuations could contaminate the phase estimate. Such contamination is possible in principle through nonlinearity in the APD or the RF receive chain, where a change in signal amplitude produces a change in output phase. This effect is known as amplitude-to-phase (AM-to-PM) conversion [Maas 2003]. The idealized model of Sec. 4.3 of the main paper predicts no such coupling. Rather than appeal to that model, we test it: we inject a large, purely optical intensity modulation while holding the scene depth fixed, so that any phase response at the injection frequency is an artifact by construction.

2.1 Signal model

We insert a liquid crystal (LC) shutter in the optical path and drive it with a sinusoidal voltage at frequency f_{sh}. The shutter acts as a time-varying transmittance that scales the returned optical power uniformly across the aperture without altering any propagation path [Batagoda et al. 2026]. Normalizing by its mean transmittance gives the envelope

$$g_{\mathrm{s}}(t) = 1 + m\cos\left(2\pi f_{\mathrm{sh}} t\right), \qquad 0 \le m \le 1, \tag{1}$$

with m the fractional modulation depth; $g_{\mathrm{s}}(t) \ge 0$ because transmittance is non-negative. The APD output is AC coupled, so the static and slowly varying terms of the received intensity are rejected and the phase-bearing photocurrent acquires $g_{\mathrm{s}}(t)$ as a multiplicative real factor,

$$i_{\mathrm{rx}}(t) = A_{\mathrm{rx}}\, g_{\mathrm{s}}(t)\cos\left(2\pi f t + \phi(t)\right). \tag{2}$$

Critically, $\phi(t)$ is untouched: the target, the optics, and all optical path lengths are static during the measurement.

Since f_{sh} (≤ 100 Hz here) lies below the intermediate frequency $f_b = 10$ kHz, the envelope passes the receiver low-pass filters intact. Quadrature demodulation and digital downconversion (Sec. 4.3) therefore return the baseband signal of the main paper scaled by that envelope,

$$s_\phi[n] = A\, g_{\mathrm{s}}[n]\, e^{j\phi[n]}, \tag{3}$$

from which the two readouts separate as

$$\left|s_\phi[n]\right| = A\, g_{\mathrm{s}}[n], \qquad \arg\left\{s_\phi[n]\right\} = \phi[n], \tag{4}$$

because multiplying a complex number by a positive real scalar scales it without rotating it. The injected modulation thus appears in full in the amplitude channel and not at all in the phase channel. Any phase component measured at f_{sh} must therefore originate in the hardware—gain compression in the APD or the SDR front-end, or an additive leakage term such as transmitter-to-receiver crosstalk—and it is this residual that the experiment below bounds.

2.2 Experimental results

Fig. 1(a) shows the SDR-ToF prototype with the LC shutter placed in the optical path.

Intensity modulation alone. With the shutter driven at 10 Hz and the speaker off (Fig. 1(b)), the shutter is enabled at approximately 2.5 s within a single continuous record. The amplitude channel responds immediately with a ten times higher peak-to-peak swing, and its spectrum shows a strong 10 Hz line. The phase channel is statistically unchanged across the transition and its spectrum contains no line at 10 Hz; its only discrete features are the 60 Hz mains line and its harmonics (spurious RF noise).

Intensity modulation with surface vibration. Adding a 500 Hz sinusoidal tone on the speaker separates the two channels cleanly (Fig. 1(c)). The amplitude spectrum contains both the 10 Hz shutter line and the 500 Hz acoustic line, since beam deflection and shutter attenuation both act on the returned intensity; the phase spectrum contains only the acoustic line. The 40 ms insets (zoomed images) resolve this directly: the amplitude trace rides on the slow shutter envelope while the phase trace carries the acoustic tone alone.

We then raise the shutter to 100 Hz (Fig. 1(d)), placing it in the same frequency decade as the acoustic tone and a factor of twenty above the 5 Hz loop bandwidth of the digital PLL (Appendix A.2 of the main paper). The result is unchanged. The amplitude spectrum shows the 100 Hz shutter line (and its harmonics), and the 500 Hz acoustic line, while the phase spectrum shows only the acoustic line and the mains harmonics.

This experiment shows that intensity and phase can carry vibration information: intensity through beam deflection, phase through path length change. Our system reads out amplitude and phase. Under the injected intensity modulation tested here, phase remains unaffected, supporting that the depth measurements reported in this paper reflect true time-of-flight rather than an intensity artifact.

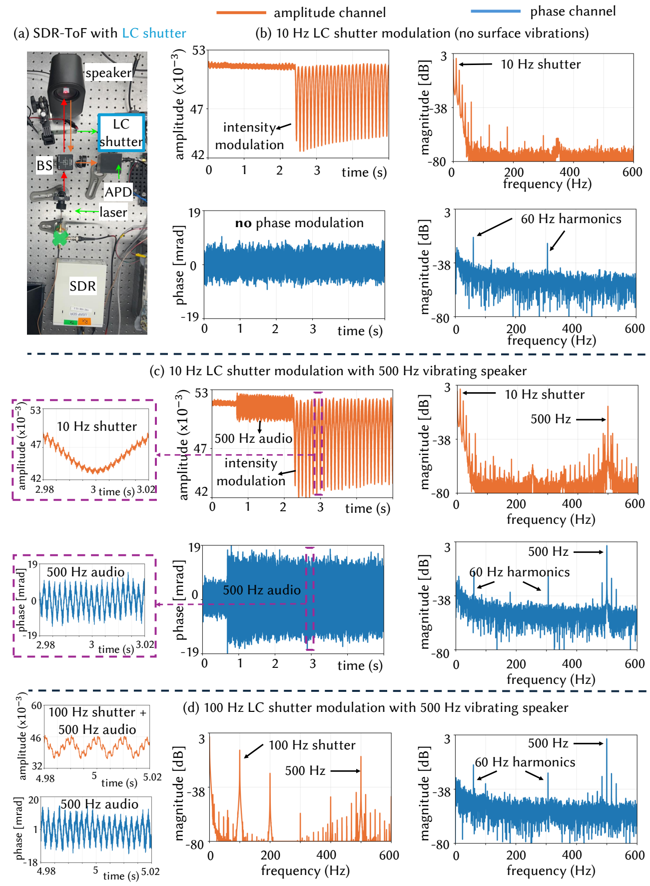

Fig. 1. **amplitude Vs phase readout under controlled optical intensity modulation.** (a) The SDR-ToF prototype with a liquid crystal (LC) shutter inserted in the optical path. The shutter modulates the returned intensity at a known frequency f_{sh} while all optical path lengths remain fixed, so any phase response at f_{sh} is an artifact. (b) Shutter at 10 Hz with the speaker off, enabled at approximately 2.5 s within a single record. The amplitude channel tracks the shutter and its spectrum shows a strong 10 Hz line. The phase channel is unchanged across the transition and shows no corresponding line; its discrete features are 60 Hz mains harmonics, present independent of the shutter. (c) Shutter at 10 Hz with a 500 Hz sinusoidal tone driving the speaker. The amplitude channel carries both the shutter modulation and the acoustic tone; the phase channel carries the acoustic tone alone. The 40 ms insets (zoomed) show the amplitude trace riding on the shutter envelope while the phase trace does not. (d) Shutter raised to 100 Hz. The amplitude spectrum shows the 100 Hz shutter line, its 200 Hz harmonic, and the 500 Hz acoustic line; the phase spectrum shows only the acoustic line and the mains harmonics. Amplitude is reported in normalized ADC full-scale units.

3 IMAGING RESOLUTION

To characterize the resolution of our depth frames empirically, we measured the modulation transfer function (MTF) using the slanted-edge method of ISO 12233 [Burns et al. 2000; ISO Standard 12233 2000]. Because a scanned system need not be isotropic, we report the two scan axes separately.

The target is a retroreflective panel with a straight edge, offset by a 5 mm step in front of a retroreflective background (Fig. 2(a–b)). The scene was raster-scanned at 100 Hz on the fast galvo axis and 1 Hz on the slow axis, sampled at 10 MS/s, giving the 200 × 250 pixel depth frame in Fig. 2(c). The measurement was then repeated with the target rotated so that the edge is slanted with respect to the slow axis (Fig. 2(e–g)). Because both surfaces are retroreflective, there is no intensity contrast at the edge, so the MTF is computed on the depth channel. The step is deliberately small: at 900 MHz a 5 mm offset corresponds to a phase difference of only 11°, so the two returns (from two surfaces) mix linearly within a boundary pixel and the measured edge is a faithful blurred replica of the true step. Larger steps introduce a nonlinearity that distorts the edge profile.

Lateral scale was calibrated independently on each axis by scanning a bar of known width (1.1 cm) along one axis at a time, at the same drive frequency and amplitude used in the 2D scan, and counting raw ADC samples across the bar. This gives 32.18 µm/px

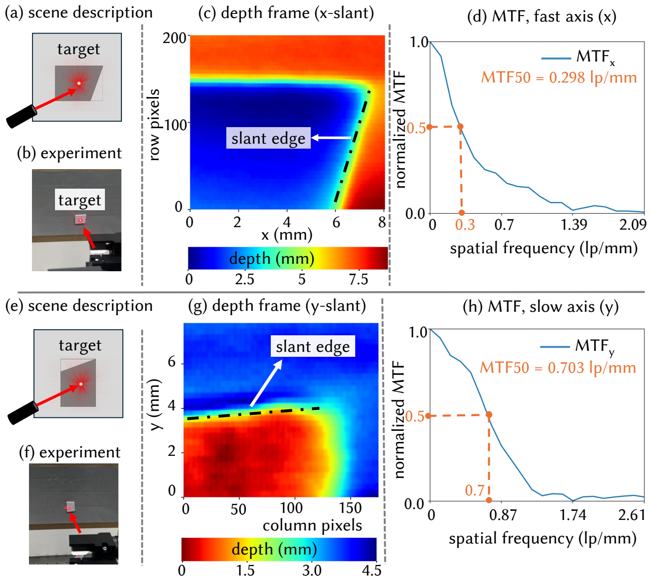

Fig. 2. **Lateral resolution measurement.** (a–d) Edge slanted with respect to the fast scan axis (x); (e–h) the same target rotated so the edge is slanted with respect to the slow scan axis (y). (a, e) Scene geometry: a retroreflective target placed 5 mm in front of a retroreflective background. (b, f) Experimental setup. (c) Reconstructed depth frame (200 × 250 px) corresponding to fast-axis slant edge. The horizontal axis is calibrated to millimeters using an independent single-axis scan (≈ 32.18 µm/px); the vertical axis is shown in row index. (g) Depth frame for the slow-axis slanted edge (150 × 175 px), with the vertical axis calibrated (≈ 52.26 µm/px) and the horizontal axis in column index. (d, h) MTF computed from the depth-channel edge spread function along the fast and slow axes respectively. MTF50 of 0.298 lp/mm (fast) and 0.703 lp/mm (slow) correspond to smallest resolved features of 1.68 mm and 0.71 mm respectively. lp—line pairs.

on the fast axis and 52.26 µm/px on the slow axis. We do not assume a shape for the blur; the slanted-edge method measures it directly. A straight edge tilted by a few degrees is crossed by successive scan lines at sub-pixel offsets; projecting the region around the edge along its fitted direction gives an oversampled edge-spread function (ESF), whose derivative is the line-spread function and whose Fourier magnitude is the MTF. Each measurement is converted to millimeters using the pitch of its own axis; projecting along that axis rather than the true edge normal is a few-percent correction and is not applied.

We measure an MTF50 of 0.298 lp/mm (lp—line pairs) along the fast scan axis (Fig. 2(d)), corresponding to a smallest resolved feature of 1.68 mm, and 0.703 lp/mm along the slow axis (Fig. 2(h)), corresponding to 0.71 mm. The slow axis is the sharper of the two because it is free of scan-induced blur: along the fast axis the beam translates appreciably during the integration window of each sample, whereas successive rows along the slow axis are separated by a full fast-axis period, so no comparable smearing occurs. The fast-axis figure is therefore the conservative bound on lateral resolving power. Both measurements characterize lateral resolving power and are distinct from depth precision, reported separately.

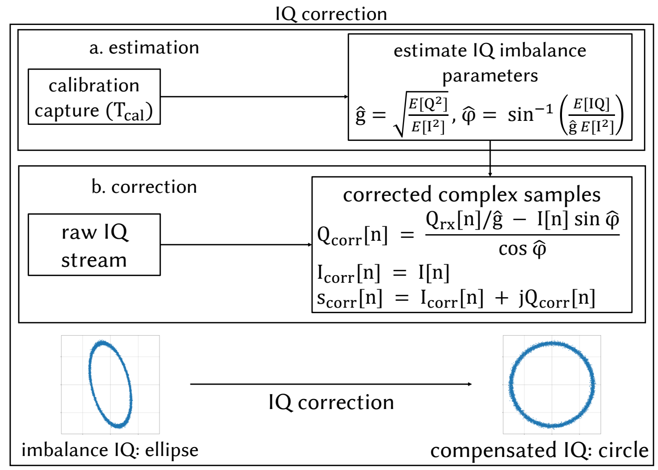

Fig. 3. IQ **imbalance estimation and correction. (a) Estimation.** A short calibration capture at the start of each acquisition is used to estimate the gain and phase imbalance parameters $\hat{g}$ and $\hat{\varphi}$ from second-order moments of the received baseband IQ samples [Windisch and Fettweis 2004]. **(b) Correction.** The estimated parameters are applied to the full raw IQ stream to invert the imbalance model and recover corrected complex samples $s_{corr}[n] = I_{corr}[n] + jQ_{corr}[n]$. The raw constellation is elliptical due to quadrature gain and phase mismatch, while the compensated samples recover the circular trajectory expected from an ideal complex sinusoid.

4 IQ IMBALANCE COMPENSATION

The downconversion process requires the local oscillator to produce sine and cosine waveforms that are precisely orthogonal and of equal amplitude. In that ideal case, the I vs Q mapping (constellation diagram) is circular. However, synthesizing quadrature waveforms at GHz carrier frequencies in silicon is a challenging task, and any mismatch in phase or amplitude between the two quadrature paths manifests as IQ imbalance in the downconverted baseband signal [Abidi 2002]. The consequence is a coupling between the I and Q channels: the received signal constellation becomes elliptical rather than circular, and an undesired image spur appears at the negative frequency of any spectral component. This imperfection produces distorted depth estimates and needs to be calibrated.

The IQ distortion model treats the in-phase channel as the undistorted reference, while the quadrature channel is subject to amplitude imbalance g and phase imbalance φ:

$$I_{rx}[n] = I[n], \tag{5}$$

$$Q_{rx}[n] = g\big(I[n]\sin\varphi + Q[n]\cos\varphi\big), \tag{6}$$

where $g = 1$ and $\varphi = 0$ correspond to the ideal, distortion-free case.

Hardware-level correction is not available on the USRP N210 SDR with WBX daughterboard; compensation is therefore performed in the digital domain. To this end, a blind estimation method based on second-order statistics is applied following the method of moments formulation of [Windisch and Fettweis 2004]. The estimator requires no pilot tones or known reference signal, operating instead on the statistical properties of the received data itself. Given a zero-mean signal (DC offset removed as a prerequisite), the amplitude and phase imbalance parameters are estimated from the empirical second-order moments $\mathbf{E}[\mathsf{I}^2]$, $\mathbf{E}[\mathsf{Q}^2]$, and $\mathbf{E}[\mathsf{IQ}]$ as

$$\hat{g} = \sqrt{\frac{\mathbf{E}[\mathsf{Q}^2]}{\mathbf{E}[\mathsf{I}^2]}}, \tag{7}$$

$$\hat{\varphi} = \sin^{-1}\left(\frac{\mathbf{E}[\mathsf{IQ}]}{\hat{g}\,\mathbf{E}[\mathsf{I}^2]}\right), \tag{8}$$

where Eq. (7) follows directly from the ratio of quadrature-to-in-phase power under the distortion model Eq. (6), and Eq. (8) is obtained by inverting the cross-moment relation $\mathbf{E}[\mathsf{IQ}] = g \sin\varphi\, \mathbf{E}[\mathsf{I}^2]$. The correction is then applied by inverting Eqs. (5) and (6):

$$\mathsf{Q}_{\mathrm{corr}}[n] = \frac{\mathsf{Q}_{\mathrm{rx}}[n]/\hat{g} - \mathsf{I}[n]\sin\hat{\varphi}}{\cos\hat{\varphi}}, \tag{9}$$

with $\mathsf{I}_{\mathrm{corr}}[n] = \mathsf{I}[n]$ unchanged. In practice, $\hat{g}$ and $\hat{\varphi}$ are estimated from a calibration segment of duration T_{cal} (tunable parameter, to less than a millisecond) at the start of each acquisition, and Eq. (9) is applied to the full signal (Fig. 3(b)). This approach is valid under the assumption that hardware parameters remain stationary over the acquisition window and we do not change SDR configuration during runtime.

5 SPECTRAL GATING FOR AUDIO RECONSTRUCTION

For audio reconstruction experiments only, we apply a custom spectral gating post-processing step inspired by the stationary noise-reduction mode of noisereduce [Sainburg 2019; Sainburg and Zorea 2025] to improve audibility (Fig. 5). Let $\Phi[f,t] = \mathrm{STFT}\{\phi_{\mathrm{audio}}[n]\}$ denote the complex STFT of the recovered phase audio, computed with $N_{\mathrm{FFT}} = 4096$ points. At a 48 kHz output rate, this gives a frequency resolution of approximately 11.7 Hz. A noise reference $\phi_{\mathrm{noise}}[n]$ is recorded separately during a quiet period, and its complex STFT is denoted by $\Phi_n[f,t] = \mathrm{STFT}\{\phi_{\mathrm{noise}}[n]\}$. We estimate per-frequency noise statistics from the magnitude of the noise STFT as

$$\mu_n[f] = \langle |\Phi_n[f,t]| \rangle_t, \qquad \sigma_n[f] = \mathrm{std}_t(|\Phi_n[f,t]|)\,.$$

A per-bin threshold is then set as

$$\tau[f] = \mu_n[f] + k\sigma_n[f],$$

with $k = 1$ in our experiments. Rather than using a hard binary gate, we form a soft sigmoid mask

$$M_{\mathrm{raw}}[f,t] = \mathrm{sigmoid}\left(\frac{|\Phi[f,t]| - \tau[f]}{\epsilon}\right),$$

where ϵ controls the softness of the transition around the threshold. The mask is smoothed with a 100 Hz kernel along the frequency axis and a 50 ms kernel along the time axis to suppress tonal gating artifacts, giving $M_{\mathrm{smooth}}[f,t]$. We then temper the attenuation using

$$M_{\mathrm{final}}[f,t] = pM_{\mathrm{smooth}}[f,t] + (1-p),$$

where $p = 0.95$ is the noise-floor attenuation factor. The denoised complex STFT is obtained by applying this mask to the original complex STFT,

$$\tilde{\Phi}[f,t] = M_{\mathrm{final}}[f,t]\Phi[f,t],$$

and the denoised time-domain phase signal is recovered as

$$\tilde{\phi}_{\mathrm{denoised}}[n] = \mathrm{ISTFT}\{\tilde{\Phi}[f,t]\}.$$

The result is shown in Fig. 4.

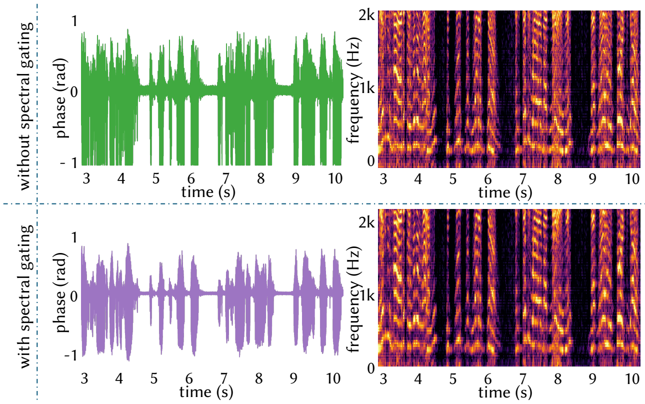

Fig. 4. **Spectral gating for phase denoising.** The top row shows the raw recovered phase signal and its corresponding spectrogram before spectral gating, while the bottom row shows the phase signal and spectrogram after spectral gating. Spectral gating suppresses broadband noise and low-energy artifacts while preserving the dominant time-varying audio/vibration structure in the recovered phase. The time-domain phase traces show reduced background fluctuations after filtering, and the spectrograms show that the main spectral components are retained after denoising.

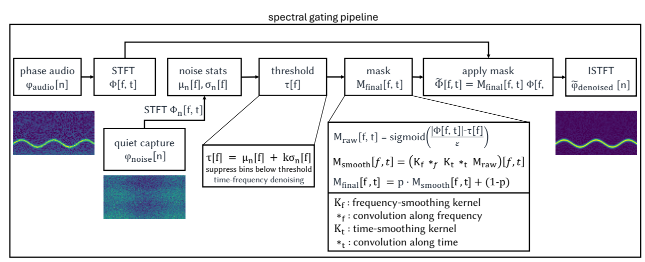

Fig. 5. **Spectral gating pipeline for audio reconstruction.** The recovered phase audio $\varphi_{\mathrm{audio}}[n]$ is transformed into a complex short-time Fourier transform (STFT) $\Phi[f,t]$. A separately recorded quiet capture $\varphi_{\mathrm{noise}}[n]$ is used to compute a noise STFT $\Phi_n[f,t]$, from which per-frequency noise statistics $\mu_n[f]$ and $\sigma_n[f]$ are estimated. These statistics define a threshold $\tau[f] = \mu_n[f] + k\sigma_n[f]$, which is used to form a soft time-frequency mask $M[f,t]$. The mask is applied to the original complex STFT, $\tilde{\Phi}[f,t] = M[f,t]\Phi[f,t]$, and the denoised phase signal $\tilde{\varphi}_{\mathrm{denoised}}[n]$ is recovered by ISTFT. Insets show the input spectrogram, the quiet-capture noise reference, and the denoised output spectrogram.

6 COST ESTIMATES FOR ALTERNATIVE HIGH-SPEED METHODS

Tab. 1 summarize the system cost estimates for Baek et al. [2023], Mirdehghan et al. [2024], Sheinin et al. [2022], visual vibrometry methods [Davis et al. 2014; Zhou et al. 2025], and SDR-ToF, respectively.

References

Asad A Abidi. 2002. Direct-conversion radio transceivers for digital communications. *IEEE Journal of solid-state circuits* 30, 12 (2002), 1399–1410.

Seung-Hwan Baek, Noah Walsh, Ilya Chugunov, Zheng Shi, and Felix Heide. 2023. Centimeter-wave free-space neural time-of-flight imaging. *ACM Transactions on Graphics* 42, 1 (2023), 1–18.

Table 1. Lower-end cost estimates for representative high-speed depth and vibration sensing systems.

Component	**Cost (USD)**
GHz correlation ToF prototype of Baek et al. [2023]	
CW Laser Source (Laser Quantum Gem 532)	$5,000
Single-Mode Fiber (OZ Optics QPMJ-A3AHPCA3AHPC-488-3.5/125)	$1,000
Resonant Electro-Optic Modulators (Custom GHz EOM, 2×)	$20,000
RF Driver + DDS + Function Generator (Siglent SDG2042X + custom RF chain)	$8,000
Polarizing Beam Splitters (Thorlabs PBS101, 2×)	$400
Half-Wave Plates (Thorlabs WPMH05M-532, 2×)	$600
Quarter-Wave Plates (Thorlabs WPQ05M-532, 2×)	$1,000
Non-Polarizing Beam Splitters (Thorlabs CCM1-BS013)	$500
Lenses / Beam Expanders (Thorlabs LC1060-A, LA1608-A, LA1213-A, LA1951-A)	$500
Mirrors (Thorlabs PF10-03-P01, PF30-03-P01 and mounts)	$1,000
2-Axis Galvo Scanner (Thorlabs GVS012)	$4,000
Avalanche Photodiode (Thorlabs APD440A)	$1,400
Reference Photodiode (Generic GHz photodiode)	$400
Low-Pass Filter (Thorlabs EF120) + RC Integrator	$300
DAQ System (LabJack T7)	$570
Integrating Sphere / Power Sensor (Thorlabs S140C)	$500
Stereo Cameras (FLIR Grasshopper3 GS3-U3-32S4C, 2×)	$2,000
Optomechanics (Breadboard, mounts, cages, rails, posts)	$5,000
RF Photodetector (EOT ET-4000, 12 GHz)	$5,000
Translation Stage (Thorlabs MTS50/M-Z8)	$2,000
Total	**$59,170**
Full-Wavefield Lidar prototype of Mirdehghan et al. [2024]	
Coherent Optical Modem (Ciena WaveLogic 5 Nano, QSFP-DD)	$10,000
Transmit EDFA Amplifier (Thorlabs EDFA300S, Erbium-Doped Fiber Amplifier)	$8,000
Receive Pre-Amplifier (Thorlabs EDFA100S, Erbium-Doped Fiber Amplifier)	$4,500
Fiber Circulator (Thorlabs 6015-3-APC, 1550 nm, 3-Port Optical Circulator)	$800
Fiber Collimator (Thorlabs CFS18-1550-APC, FiberPort Collimator)	$500
2-Axis Galvo Scanning Mirrors (Thorlabs GVS012, Galvanometer Scanner)	$4,000
DAQ / Mirror Controller (National Instruments USB-6341, Multifunction I/O Device)	$1,000
Host PC + QSFP-DD Interface (High-Speed Data Interface + Control Workstation)	$2,000
Single-Mode Fiber Patch Cables (FC/APC, SMF-28, multiple lengths)	$600
Fiber Attenuators / Connectors (FC/APC adapters, inline attenuators)	$400
Optical Breadboard (Thorlabs MB3045) + Posts, Mounts, and Clamps	$2,000
Power Supplies and Control Electronics (DC supplies, drivers, synchronization)	$1,000
Laser Safety Equipment (1550 nm safety goggles, beam blocks, IR cards)	$1,000
Total	**$35,800**
Prototype of Sheinin et al. [2022]	
Rolling-Shutter Camera (IDS UI3240-NIR-GL)	$580
Global-Shutter Camera (IDS UI3070CP-M-GL)	$1,000
Relay Optics	$700
Objective Lens (16 mm, f/1.8)	$250
Cylindrical Lens	$100
532 nm Laser	$200
Diffraction Gratings	$150
Beam Splitter for Coaxial Laser Path	$250
Total	**$3,230**
Visual vibrometry of Davis et al. [2015] and Zhou et al. [2025]	
High-Speed Camera (Phantom v10, for Davis et al. [2015])	$3,500
Event Camera (Prophesee EVK4 or similar, for Zhou et al. [2025])	$3,500
Total	**$3,500 (either camera)**
SDR-ToF prototype (ours)	
Full-Duplex SDR Transceiver (USRP N210)	$2,500
Laser Diode (Thorlabs LP638P150, 638 nm, 150 mW)	$70
Laser Diode Mount (Thorlabs S05LM38)	$35
DC Power Supply (Amazon)	$50
Wideband Bias-Tee (Mini-Circuits ZFBT-4R2G+)	$70
2-Axis Galvo Scanner (Thorlabs GVS012)	$4,000
Focusing Lens (Thorlabs LA1612-A)	$30
Avalanche Photodetector (Menlo Systems APD210)	$2,800
Beam splitter (Thorlabs BS013)	$260
Total	**$9,815**

Nevindu M Batagoda, Ramchander Rao Bhaskara, Christopher Metzler, and Adithya Pediredla. 2026. ShutterEvents: Shutter Modulation for Static-Dynamic Scene Recovery. (2026).

Peter D Burns et al. 2000. Slanted-edge MTF for digital camera and scanner analysis. In *Is and Ts Pics Conference*. Society for Imaging Science & Technology, 135–138.

Abe Davis, Katherine L Bouman, Justin G Chen, Michael Rubinstein, Fredo Durand, and William T Freeman. 2015. Visual vibrometry: Estimating material properties from small motion in video. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*. 5335–5343.

Abe Davis, Michael Rubinstein, Neal Wadhwa, Gautham J Mysore, Frédo Durand, and William T Freeman. 2014. The visual microphone: passive recovery of sound from video. *ACM Transactions on Graphics (TOG)* 33, 4 (2014), 1–10.

ISO Standard 12233. 2000. Photography—Electronic Still-Picture Cameras—Resolution Measurements.

Stephen A Maas. 2003. *Nonlinear microwave and RF circuits.* Artech house.

Parsa Mirdehghan, Brandon Buscaino, Maxx Wu, Doug Charlton, Mohammad E Mousa-Pasandi, Kiriakos N Kutulakos, and David B Lindell. 2024. Coherent optical modems for full-wavefield lidar. In *SIGGRAPH Asia 2024 Conference Papers.* 1–10.

Tim Sainburg. 2019. *timsainb/noisereduce: v1.0.* doi:10.5281/zenodo.3243139

Tim Sainburg and Amnon Zorea. 2025. Domain general noise reduction for time series signals with Noisereduce. *Scientific Reports* 15 (Aug. 2025), 30905. doi:10.1038/s41598-025-13108-x

Mark Sheinin, Dorian Chan, Matthew O'Toole, and Srinivasa G Narasimhan. 2022. Dual-shutter optical vibration sensing. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition.* 16324–16333.

Marcus Windisch and Gerhard Fettweis. 2004. Blind I/Q imbalance parameter estimation and compensation in low-IF receivers. In *First International Symposium on Control, Communications and Signal Processing, 2004.* IEEE, 75–78.

Xinyu Zhou, Peiqi Duan, Yeliduosi Xiaokaiti, Chao Xu, and Boxin Shi. 2025. Event-based Visual Vibrometry. In *Proceedings of the IEEE/CVF International Conference on Computer Vision.* 24666–24676.