A universal 3D imaging sensor on a silicon photonics platform

Christopher Rogers[1,*], Alexander Y. Piggott[1,*], David J. Thomson[2], Robert F. Wiser[1], Ion E. Opris[3], Steven A. Fortune[1], Andrew J. Compston[1], Alexander Gondarenko[1], Fanfan Meng[2], Xia Chen[2], Graham T. Reed[2], and Remus Nicolaescu[1,†]

[1]*Pointcloud Inc, San Francisco, California 94107, USA*

[2]*Optoelectronics Research Centre, University of Southampton, University Road, Southampton, Hampshire SO17 1BJ, UK*

[3]*Opris Consulting, San Jose, California 95138, USA*

∗*These authors contributed equally to the work.*

†*Corresponding author: remus.nicolaescu@point.cloud*

This paper is dedicated to the memory of Sunil Sandhu.

Accurate 3D imaging is essential for machines to map and interact with the physical world[1,2]. While numerous 3D imaging technologies exist, each addressing niche applications with varying degrees of success, none have achieved the breadth of applicability and impact that digital image sensors have achieved in the 2D imaging world[3–10]. A large-scale two-dimensional array of coherent detector pixels operating as a light detection and ranging (LiDAR) system could serve as a universal 3D imaging platform. Such a system would offer high depth accuracy and immunity to interference from sunlight, as well as the ability to directly measure the velocity of moving objects[11]. However, due to difficulties in providing electrical and

photonic connections to every pixel, previous systems have been restricted to fewer than 20 pixels[12–15]. Here, we demonstrate the first large-scale coherent detector array consisting of 512 (32×16) pixels, and its operation in a 3D imaging system. Leveraging recent advances in the monolithic integration of photonic and electronic circuits, a dense array of optical heterodyne detectors is combined with an integrated electronic readout architecture, enabling straightforward scaling to arbitrarily large arrays. Meanwhile, two-axis solid-state beam steering eliminates any tradeoff between field of view and range. Operating at the quantum noise limit[16,17], our system achieves an accuracy of $3.1\ \mathrm{mm}$ at a distance of 75 metres using only $4\ \mathrm{mW}$ of light, an order of magnitude more accurate than existing solid-state systems at such ranges. Future reductions of pixel size using state-of-the-art components could yield resolutions in excess of 20 megapixels for arrays the size of a consumer camera sensor. This result paves the way for the development and proliferation of low cost, compact, and high performance 3D imaging cameras, enabling new applications from robotics and autonomous navigation to augmented reality and healthcare.

The digital complementary metal–oxide–semiconductor (CMOS) image sensor revolutionized 2D imaging, borrowing technology from silicon microelectronics to produce a flexible and scalable camera sensor[18]. As a focal plane array (FPA), the digital image sensor operates in concert with a lens that focuses light and forms an image on the detector. A key advantage of this scheme is that the field of view and light collection efficiency are not set by the image sensor, but instead by the choice of lens. Furthermore, the CMOS image sensor can be optimized for high performance or cost, allowing it to be fine-tuned for different applications. Due to the great flexibility afforded by this arrangement, the digital CMOS sensor has become the sensor of choice for the majority of 2D imaging.

In contrast, the world of 3D imaging is characterized by a vast assortment of competing technologies, each addressing a small niche of applications. Long range and high precision applications such as autonomous vehicles and construction site mapping are dominated by expensive and fragile mechanically steered light detection and ranging (LiDAR) systems[3,4]. Meanwhile, solid-state solutions such as structured light[5] and time-of-flight arrays[6–10,19,20] are used when affordability, compactness, and reliability must be achieved at the expense of performance, such as in mobile devices and augmented reality systems. Optical phased arrays are a promising solid-state approach, but the development of long-range 2D-scanning systems has proven challenging, with current demonstrations limited to less than 20 metres[21–23]. As such, no currently available technology can address the needs of these diverse use cases.

Here, we demonstrate a fully solid-state, integrated photonic LiDAR based on the same FPA concept as the CMOS image sensor. By making efficient use of light, our system achieves the long range and high depth accuracy needed for demanding applications such as self-driving vehicles[1] and drone-based 3D mapping[24,25]. The architecture also scales to arbitrarily large fields of view. The centerpiece of our system is the coherent receiver array, a highly sensitive array of compact optical heterodyne detectors operating at the quantum noise limit[16,17]. To eliminate any tradeoff between field of view and range, the receiver is paired with a solid-state beam steering mechanism that sequentially illuminates the scene in small patches.

The coherent receiver array allows our architecture to operate using the robust frequency-modulated continuous-wave (FMCW) coherent LiDAR scheme[26,27]. In contrast to widely used time-of-flight LiDARs that rely on transmitting short pulses of light, an FMCW LiDAR uses a linearly chirped laser as both the transmit beam and the local oscillator. Scattered light received from the target is mixed with the local oscillator light in a heterodyne receiver, producing a beat frequency proportional to the round trip travel time, and hence the distance to the target. Due to the use of coherent detection, the receiver detects a single polarization of scattered light. In coherent LiDAR systems, the receiver polarization is typically chosen to be the same as the transmitter polarization[11,26,27] as is done here, since most materials preferentially scatter light into the same polarization as the illuminating light[28].

The FMCW scheme confers our architecture with a number of key advantages relative to time-of-flight schemes. First, due to the use of heterodyne detection, the system is immune to interference from sunlight and other LiDAR systems operating nearby since it selectively detects light close in frequency to the local oscillator light[11]. Second, a coherent LiDAR can directly measure the velocity of moving objects by sensing the Doppler shift of the received signal[26,27]. Third, high depth accuracy is straightforward to achieve since it depends upon only the chirp bandwidth and signal-to-noise ratio[29], allowing the receiver electronics to operate at relatively low frequencies. This is in contrast to time-of-flight schemes where depth accuracy is limited by receiver bandwidth. Finally, the FMCW system is well suited for integrated photonic LiDARs, which are constrained in peak power due to nonlinear effects[30,31]. Whereas time-of-flight schemes emit photons in short high-power bursts, the FMCW scheme emits photons continuously and maximizes the number of emitted photons, thereby improving the system's range.

Despite the numerous advantages of a 3D imaging system based on the coherent receiver array, previous demonstrations have been limited to fewer than 20 pixels due to their reliance on direct electrical connections to each pixel[12–15]. To solve this issue of scalability, we implemented our LiDAR system on GlobalFoundries' CMS90WG process, a silicon photonics process with monolithically integrated CMOS electronics[32]. This allowed us to incorporate a highly multiplexed electronic readout architecture directly into the receiver array, minimizing the number of required external electrical connections while maintaining signal integrity. Our prototype array contains 512 pixels, and can be scaled to arbitrarily large numbers of pixels by simply increasing the size of the array. Furthermore, due to the use of a standard process provided by a commercial foundry, our system can immediately be mass produced for minimal cost.

Scalable 3D imaging architecture

As shown in Fig. 1(a), our architecture is based on two FPAs. The first acts as a transmitter, and the second as a receiver. Chirped laser light for the FMCW scheme is generated external to the chip by modulating a fixed-frequency $1550\ \mathrm{nm}$ laser with a silicon-photonic IQ Mach-Zehnder modulator (MZM), which is in turn driven by an arbitrary waveform generator. This approach ensures chirp linearity and enables the use of a simple, low-noise laser.

Long-range performance is achieved by sequentially illuminating and reading out the scene in small patches. By only illuminating pixels that are currently being read out, light is used as efficiently as possible. As illustrated in Fig. 1(b), this is accomplished on the transmitter side by a switching tree terminated by a FPA of grating couplers. Light is directed to one transmit grating at a time, illuminating a small subset of the scene. This switching approach to beam steering is robust and can be scaled up to arbitrarily large arrays, with optical losses limited only by waveguide scattering[33] and the extinction ratio of the switching trees. Meanwhile, the receiver consists of a dense FPA of miniaturized heterodyne receivers. All receiver pixels that correspond to the illuminated area are simultaneously read out in parallel. Since the angular resolution is defined by the point spread function of the lens, which drops off very quickly, there is negligible crosstalk between different receiver pixels. To avoid wasting local oscillator light, a second switching tree is used to provide only the activated subset of the receiver FPA with local oscillator light.

The use of parallel readout in the receiver is fundamental to the scalability of our architecture. First, the system resolution is defined by the number of pixels in the receiver FPA, rather than the number of steering positions. This significantly improves the system resolution for a given chip size since heterodyne receiver pixels are roughly an order of magnitude smaller than thermo-optic switches. Second, parallel readout eliminates the need for fast thermo-optic switching because the number of measured points per second is decoupled from the switching rate. Finally, due to the use of an FMCW scheme, parallel readout proportionally reduces the receiver signal frequencies by allowing longer ramp times, simplifying the readout electronics.

Implementation on a hybrid CMOS-photonics process

An optical micrograph of our demonstrator chip is shown in Fig. 1(c). The transmitter consists of a 1×16 thermo-optic switch tree with 16 grating couplers in the transmit FPA. Meanwhile, the receiver consists of a 32×16 (512) pixel array of heterodyne receivers, with local oscillator light provided by a 1×8 switch tree. Operation of the thermo-optic switching trees is demonstrated in Extended Data Fig. 1. The thermo-optic switching trees were found to be very stable, with no recalibration required even after several months of operation in an uncontrolled temperature environment.

As schematically illustrated in Fig. 2(a), the receiver FPA consists of a tiled array of miniaturized heterodyne receiver pixels. Each pixel collects scattered light from the scene using a grating coupler. Meanwhile, local oscillator light is provided to each pixel via a network of silicon waveguides. The scattered light and local oscillator (LO) are mixed on a balanced detector consisting of a 50-50 directional coupler and germanium PIN photodiodes, producing a heterodyne tone in the electrical domain corresponding to the target's distance. The signal is then amplified by a transimpedance amplifier (TIA) integrated within the pixel. A buffer amplifier at the end of each row of pixels is shared among the pixels of that row, and maintains wide bandwidth while driving the large parasitic capacitances of the wiring and multiplexed circuitry. Simultaneously active driver amplifiers carry the signal to the edge of the chip, enabling parallel readout. As shown in Fig. 2(b), the individual pixels are turned on and off using a power switch built into each TIA, and an inter-stage RC filter flattens the frequency response to simplify downstream signal processing.

In general, minimizing the input-referred noise of the electronic signal chain improves the pixel's sensitivity and detection probability. Furthermore, higher receiver bandwidths are desirable when using the FMCW scheme since this reduces the required integration time for a given maximum range. In our architecture, the TIA feedback resistance determines the gain, bandwidth, and noise, with bandwidth and noise decreasing with larger resistance[34]. Due to the use of compact waveguide-coupled photodiodes and tight integration between the photodiodes and TIAs, we have a remarkably small parasitic capacitance of only 1.5 fF. As a result, we achieve very low noise performance in the electrical signal chain with $20\ \mathrm{k\Omega}$ of gain and bandwidths above $280\ \mathrm{MHz}$, as shown in Fig. 3(a)-(b). Note that this is the measured bandwidth through the packaged chip; the simulated bandwidth of the on-chip amplifier chain is $750\ \mathrm{MHz}$, suggesting a bandwidth limitation in the test setup. Even so, as seen in Fig. 3(d), our integrated TIA design allows similar gain-bandwidth product with $2-3\times$ lower noise floor as compared to published conventional systems, where the photodiodes and the amplifier chains are on separate chips.

Due to the low noise of our on-chip amplifier chain, our receiver FPA operates at the quantum limit for sensitivity, which is reached when local oscillator shot noise dominates all other noise sources[16,17]. As shown in Fig. 3(e), shot noise reaches parity with amplifier noise with only $5\ \mu\mathrm{W}$ of LO power for a typical pixel in the receiver array, in contrast to coherent receivers used for telecommunications applications which typically require one to two orders of magnitude more LO power. Combined with the excellent $30-40\ \mathrm{dB}$ common-mode rejection ratio of the balanced heterodyne detectors, as shown in Fig. 3(c), this makes the receiver array significantly less susceptible to local oscillator noise sources such as laser relative intensity noise, optical amplifier noise, and chirp generator noise. Furthermore, the low LO power per pixel significantly reduces the number of required thermo-optic switches for LO distribution since many receiver pixels can simultaneously share LO power.

Monolithic integration of electronics into the receiver FPA facilitates the use of an actively multiplexed readout architecture, allowing the receiver to be scaled to arbitrarily large numbers of pixels. In our demonstrator chip, multiple levels of multiplexing and amplification are used to map 512 pixels to 8 outputs while maintaining signal integrity. As illustrated in Extended Data Fig. 2(a-c), the pixels are read out in blocks of 8 at a time. The lowest level of multiplexing is achieved by making use of the power switch incorporated into each pixel's TIA: only one pixel per row is activated at a time. The appropriate receiver block is then selected by activating the set of eight buffer amplifiers associated with that block. A final set of differential output amplifiers drives eight off-chip $100\ \Omega$ loads. The output analog signals are fed into a bank of off-chip analog-to-digital converters for digitization, followed by digital signal processing on a field-programmable gate-array (FPGA).

The transmitter illumination pattern is closely synchronized with the receiver readout pattern, as detailed in Extended Data Fig. 2(d). Ideally, the transmitter illumination pattern should exactly match the readout pattern, so that only the receiver pixels currently being read out are illuminated by the transmitter. However, in our current prototype, each transmitter steering position illuminates the field of view of 32 receiver pixels, and 8 receiver pixels are read out at a time. This mismatch was due to a combination of chip area constraints and particularly large $1\ \mathrm{mm}$ long thermo-optic phase shifters, and can be resolved by using existing designs for compact thermo-optic shifters [35]. To further improve the optical efficiency of the system[36], a microlens array was placed in front of the transmitter array as illustrated in Extended Data Fig. 3. This produced a structured illumination pattern that exactly matched the grating coupler positions in the receiver FPA, as shown in Extended Data Fig. 4, yielding a $24\times$ improvement in signal strength.

3D imaging and velocimetry

Operation of the full LiDAR system is presented in Fig. 4. Our 3D imager was operated with an emitter power of only $4\ \mathrm{mW}$ at the aperture, a chirp bandwidth of $4\ \mathrm{GHz}$, and up- and down-chirp lengths of $850\ \mu\mathrm{s}$. As shown in Fig. 4(a), distance and velocity are encoded in the frequencies of the tones detected by each pixel[26,27]. As demonstrated in Fig. 4(b), the system achieved a measurement precision of $1.8\ \mathrm{mm}$ at $17\ \mathrm{m}$ for a 85% reflectance target, and $3.1\ \mathrm{mm}$ at $75\ \mathrm{m}$ for a 30% reflectance target, as detailed in the Methods. Due to the effects of speckle, which equally impacts all coherent radar and LiDAR schemes[36], the detection probability was 97% for the $17\ \mathrm{m}$ target, and 42% for the $75\ \mathrm{m}$ target. Meanwhile, the measured velocity precision for slowly moving objects is $1.02\ \mathrm{mm/s}$, as shown in Fig. 4(c). Point clouds of a rotating basketball at $17\ \mathrm{m}$, stacked boxes at $55\ \mathrm{m}$, and an exterior wall are illustrated in Fig. 4(d-h). The point clouds were generated by stacking 3 sequential frames to minimize the number of missing pixels due to speckle effects. No incoherent averaging was used. Detection probabilities were calculated without any frame stacking. The missing band of points in the middle of the point clouds is due to a narrow gap in the receiver array for electrical and optical routing, as shown in Fig. 1(c). In future designs, this gap can easily be reduced or eliminated by more aggressive chip layout and routing.

Discussion and outlook

We have demonstrated a scalable solid-state 3D imaging architecture that achieves > 70 m range and millimetre-class accuracy, all while using only 4 mW of transmitted power. Correcting for the $4\times$ mismatch between the number of transmitter and receiver positions discussed earlier, this is equivalent to an optical efficiency of $0.2\ \mu\mathrm{J/point}$. Our 3.1 mm precision is an order of magnitude higher than existing solid-state 3D imagers at these ranges, with state-of-the-art flash LiDAR systems limited to an accuracy of several centimetres for distances greater than 50 metres[7–10]. Furthermore, this level of performance meets the needs of a variety of demanding applications that were previously out of reach for solid-state 3D imaging systems. For example, self-driving vehicles need a LiDAR that uses low levels of laser energy to remain eye safe, but can still achieve long ranges and high accuracy [1]. Currently, this combination of requirements is typically met using mechanically steered LiDARs, such as the commonly used Velodyne VLP-16. This 100 m class mechanical LiDAR uses the same $0.2\ \mu$J of light per point as our system[4], and has a much poorer depth accuracy of 3 cm. Meanwhile, the 3D mapping of buildings and construction sites using drones[24,25] and stationary scanners[2] requires millimetre-class accuracy at distances of tens of metres[37], which is easily achieved by our system.

System range could be further improved through increases in transmitter power, which can be readily achieved by optimizing the silicon photonic elements to minimize transmission losses.

Reducing the effects of two-photon absorption with larger waveguides or reverse-biased PN junctions would lead to further increases in transmitter power, with previous demonstrations reaching optical powers on the order of 1 W [31]. Since the range of a coherent LiDAR scales as the square root of transmitter power[36], this implies that our architecture could operate at ranges of up to 1 km. Conversely, for a fixed distance of 75 m increasing the transmitter power to 1 W would increase our system's point rate from 4.7×10^3 to 5×10^6 points/s. The current depth accuracy of $2-3$ mm could also be improved by increasing the chirp bandwidth of 4 GHz. Demonstrations of 50 GHz silicon photonic modulators [38] imply that depth accuracies of ~ 200 μm are feasible.

Our 3D imaging architecture naturally scales to large arrays. On the transmitter side, due to the use of a $1 \times N$ switching tree to steer light, the number of active thermo-optic switches, and therefore the electrical power consumption, scales as $O(\log N)$. Consequently, the number of DACs and drivers needed to control the thermo-optic switching trees also scales as $O(\log N)$. Meanwhile, on the receiver side, only the active amplifiers in the receiver chain are powered on at any given time. The power consumption of the receiver electronics thus only depends upon the number of parallel readout channels, and is essentially independent of the size of the array. Finally, due to the use of on-chip multiplexing made possible by the use of monolithically integrated electronics, the number of external electrical connections needed to interface with the chip scales as $O(\log N)$.

The fundamental limit on array size therefore comes from the size of the chip. The switching trees are negligible in size compared to the receiver FPA if existing compact designs for thermo-optic phase shifters are used. Efficient thermo-optic phase shifters as short as $35\ \mu\mathrm{m}$ in length have previously been demonstrated [35]. At the current receiver pixel pitch of $80 \times 100\ \mu\mathrm{m}^2$, chips the size of a full-frame camera sensor ($36 \times 24\ \mathrm{mm}^2$) would therefore correspond to QVGA (320×240 pixel) resolution. However, the current pixel size is limited by the use of foundry PDK devices, which were not designed to minimize footprint. Using state-of-the-art designs, $8 \times 5\ \mu\mathrm{m}^2$ pixels are feasible. Photodiodes with a footprint of $3 \times 1\ \mu\mathrm{m}^2$ are enabled by the short absorption length of germanium[39]. Meanwhile, efficient grating couplers[40] with a footprint of $3 \times 3\ \mu\mathrm{m}^2$, and 2×2 couplers[41] as small as $3 \times 1\ \mu\mathrm{m}^2$ have been demonstrated. Employing such designs, a full-frame sensor with 4500×4800 pixel resolution could be readily achieved, and further design and process refinements should yield even higher resolutions.

In conclusion, we have developed a universal solid-state 3D imaging architecture that has the potential to meet the needs of nearly all 3D imaging applications, spanning from robotics and autonomous navigation to consumer products such as augmented reality headsets. Our results suggest that the equivalent of the CMOS image sensor for 3D imaging is imminent, ushering in a broad range of applications which were previously impractical or unimaginable.

1. Urmson, C. *et al.* Autonomous driving in urban environments: Boss and the Urban Challenge. *Journal of Field Robotics* **25**, 425–466 (2008).

2. Wang, Q. & Kim, M.-K. Applications of 3D point cloud data in the construction industry: A fifteen-year review from 2004 to 2018. *Advanced Engineering Informatics* **39**, 306 – 319 (2019).

3. Lichti, D. D. Error modelling, calibration and analysis of an AM-CW terrestrial laser scanner system. *ISPRS Journal of Photogrammetry and Remote Sensing* **61**, 307–324 (2007).

4. Kidd, J. *Performance Evaluation of the Velodyne VLP-16 System for Surface Feature Surveying*. Master's thesis, University of New Hampshire (2017). URL `https://scholars.unh.edu/thesis/1116`.

5. Salvi, J., Pagès, J. & Batlle, J. Pattern codification strategies in structured light systems. *Pattern Recognition* **37**, 827 – 849 (2004). Agent Based Computer Vision.

6. Corti, A., Giancola, S., Mainetti, G. & Sala, R. A metrological characterization of the Kinect V2 time-of-flight camera. *Robotics and Autonomous Systems* **75**, 584–594 (2016).

7. McManamon, P. Review of LADAR: a historic, yet emerging, sensor technology with rich phenomenology. *Optical Engineering* **51**, 1 – 14 (2012).

8. McManamon, P. F. *et al.* Comparison of flash LiDAR detector options. *Optical Engineering* **56**, 1 – 23 (2017).

9. Hutchings, S. W. *et al.* A reconfigurable 3-D-stacked SPAD imager with in-pixel histogramming for flash LIDAR or high-speed time-of-flight imaging. *IEEE Journal of Solid-State Circuits* **54**, 2947–2956 (2019).

10. Ronchini Ximenes, A. *et al.* A modular, direct time-of-flight depth sensor in 45/65-nm 3-D-stacked CMOS technology. *IEEE Journal of Solid-State Circuits* **54**, 3203–3214 (2019).

11. Behroozpour, B., Sandborn, P. A. M., Wu, M. C. & Boser, B. E. Lidar system architectures and circuits. *IEEE Communications Magazine* **55**, 135–142 (2017).

12. Aflatouni, F., Abiri, B., Rekhi, A. & Hajimiri, A. Nanophotonic coherent imager. *Opt. Express* **23**, 5117–5125 (2015).

13. Martin, A. *et al.* Photonic integrated circuit-based FMCW coherent LiDAR. *Journal of Lightwave Technology* **36**, 4640–4645 (2018).

14. Inoue, D., Ichikawa, T., Kawasaki, A. & Yamashita, T. Demonstration of a new optical scanner using silicon photonics integrated circuit. *Opt. Express* **27**, 2499–2508 (2019).

15. Li, C., Cao, X., Wu, K., Li, X. & Chen, J. Lens-based integrated 2d beam-steering device with defocusing approach and broadband pulse operation for lidar application. *Opt. Express* **27**, 32970–32983 (2019).

16. Collett, M., Loudon, R. & Gardiner, C. Quantum theory of optical homodyne and heterodyne detection. *Journal of Modern Optics* **34**, 881–902 (1987).

17. Rubin, M. A. & Kaushik, S. Squeezing the local oscillator does not improve signal-to-noise ratio in heterodyne laser radar. *Opt. Lett.* **32**, 1369–1371 (2007).

18. El Gamal, A. & Eltoukhy, H. Cmos image sensors. *IEEE Circuits and Devices Magazine* **21**, 6–20 (2005).

19. Stann, B. L. *et al.* A 32x32 pixel focal plane array ladar system using chirped amplitude modulation. In Kamerman, G. W. & Kamerman, G. W. (eds.) *Laser Radar Technology and*

Applications IX, vol. 5412, 264 – 272. International Society for Optics and Photonics (SPIE, 2004).

20. Hu, K. *et al.* Design of a CMOS ROIC for InGaAs self-mixing detectors used in FM/cw LADAR. *IEEE Sensors Journal* **17**, 5547–5557 (2017).

21. Poulton, C. V. *et al.* Coherent solid-state LIDAR with silicon photonic optical phased arrays. *Opt. Lett.* **42**, 4091–4094 (2017).

22. Miller, S. A. *et al.* 512-element actively steered silicon phased array for low-power LIDAR. In *Conference on Lasers and Electro-Optics*, JTh5C.2 (Optical Society of America, 2018).

23. Poulton, C. V. *et al.* Long-range LiDAR and free-space data communication with high-performance optical phased arrays. *IEEE Journal of Selected Topics in Quantum Electronics* **25**, 1–8 (2019).

24. Wang, J. *et al.* Integrating BIM and LiDAR for real-time construction quality control. *Journal of Intelligent & Robotic Systems* **79**, 417–432 (2015).

25. Kasturi, A., Milanovic, V., Atwood, B. H. & Yang, J. UAV-borne lidar with MEMS mirror-based scanning capability. In Turner, M. D. & Kamerman, G. W. (eds.) *Laser Radar Technology and Applications XXI*, vol. 9832, 206 – 215. International Society for Optics and Photonics (SPIE, 2016).

26. Griffiths, H. D. New ideas in FM radar. *Electronics Communication Engineering Journal* **2**, 185–194 (1990).

27. Riemensberger, J. *et al.* Massively parallel coherent laser ranging using a soliton microcomb. *Nature* **581**, 164–170 (2020).

28. Chen, H.-S. & Rao, C. R. N. Polarization of light on reflection by some natural surfaces. *Journal of Physics D: Applied Physics* **1**, 1191–1200 (1968).

29. Thurn, K., Ebelt, R. & Vossiek, M. Noise in homodyne FMCW radar systems and its effects on ranging precision. In *2013 IEEE MTT-S International Microwave Symposium Digest (MTT)*, 1–3 (2013).

30. Tsang, H. K. *et al.* Optical dispersion, two-photon absorption and self-phase modulation in silicon waveguides at 1.5μm wavelength. *Applied Physics Letters* **80**, 416–418 (2002).

31. Rong, H. *et al.* An all-silicon Raman laser. *Nature* **433**, 292–294 (2005).

32. Giewont, K. *et al.* 300-mm monolithic silicon photonics foundry technology. *IEEE Journal of Selected Topics in Quantum Electronics* **25**, 1–11 (2019).

33. Shen, Y. *et al.* Deep learning with coherent nanophotonic circuits. *Nature Photonics* **11**, 441–446 (2017).

34. Razavi, B. The transimpedance amplifier [a circuit for all seasons]. *Solid-State Circuits Magazine* **11**, 10–97 (2019).

35. Mendez-Astudillo, M., Okamoto, M., Ito, Y. & Kita, T. Compact thermo-optic MZI switch in silicon-on-insulator using direct carrier injection. *Opt. Express* **27**, 899–906 (2019).

36. Wang, J. Y. Heterodyne laser radar-SNR from a diffuse target containing multiple glints. *Appl. Opt.* **21**, 464–476 (1982).

37. Rebolj, D., Pučko, Z., Babič, N. Čuš., Bizjak, M. & Mongus, D. Point cloud quality requirements for scan-vs-BIM based automated construction progress monitoring. *Automation in Construction* **84**, 323 – 334 (2017).

38. Samani, A. *et al.* Silicon photonic Mach–Zehnder modulator architectures for on chip PAM-4 signal generation. *Journal of Lightwave Technology* **37**, 2989–2999 (2019).

39. Sorianello, V. *et al.* Near-infrared absorption of germanium thin films on silicon. *Applied Physics Letters* **93**, 111115 (2008).

40. Sun, J., Timurdogan, E., Yaacobi, A., Hosseini, E. S. & Watts, M. R. Large-scale nanophotonic phased array. *Nature* **493**, 195–199 (2013).

41. Piggott, A. Y. *et al.* Inverse-designed photonics for semiconductor foundries. *ACS Photonics* **7**, 569–575 (2020).

Methods

Design and fabrication. The demonstration chips used as transmitter and receiver FPAs were fabricated using GlobalFoundries' CMS90WG 300 mm silicon photonics process, which monolithically integrates photonic devices with $90\ \mathrm{nm}$ silicon-on-insulator (SOI) radio-frequency (RF) CMOS electronics. All photonic devices used in the design, with the exception of the directional couplers, were provided in the foundry's standard process development kit (PDK). By doing so the photonic architecture had correct-by-construction device placement and connectivity, verifiable using Mentor Graphics' Calibre Design Rule Checker. The integrated electronics followed a standard design flow using Cadence Virtuoso and Spectre for circuit design and layout, and Mentor Graphics' Calibre for verification of design rules, comparing layout-versus-schematic, and extracting parasitics. The two domains are merged into a single hierarchy enabling connectivity verification at the receiver photodiodes along with design rule verification of closely intertwined photonics and electronics across the chip.

Due to the limited number of dies available from the multi-project wafer shuttle, we were able to fully test the functionality of 5 dies. We did not observe any defects such as dead pixels or inoperative thermo-optic switches across these dies.

Optical chirp scheme. A linearly chirped optical field $E(t)$ has the form

$$E(t) = \exp\left(i2\pi f_0 t + i\pi r t^2\right) = \left[\cos\pi r t^2 + i\sin\pi r t^2\right]\exp(i2\pi f_0 t), \tag{1}$$

where f_0 is the carrier frequency, and r is the chirp ramp rate. Thus, by coherently modulating fixed-frequency light with a microwave chirp of the form $\cos\pi r t^2 + i\sin\pi r t^2$, we produce a linear chirp in the optical domain.

In our demonstrator system, we use a series of linear up-chirps immediately followed by a series of down-chirps. The mean and difference of the up-chirp and down-chirp beat frequencies allow separate measurement of range and velocity of a target respectively[11,26,27].

For low-velocity measurements, we use a single up-chirp and single down-chirp, each with a length of $850\ \mu$s, as illustrated in Extended Data Figure 5(c). The beat frequencies for each chirp were then extracted using fast-Fourier transforms (FFTs). For fast moving objects with sufficiently large Doppler shifts, however, the beat frequency wraps around zero, resulting in ambiguous measurements. To compensate for this effect, high velocity measurements were performed using a series of fifty $17\ \mu$s up-chirps, followed by fifty $17\ \mu$s down-chirps, as shown in Extended Data Figure 5(d). The multiple chirps were coherently combined using a two-dimensional FFT to extract the beat frequencies, maintaining the same signal-to-noise ratio as the single-chirp case [m1,m2].

Optical setup. A narrow-linewidth (< 100 Hz) fiber laser (NKT Adjustik) operating at 1550 nm was used as the seed laser for the FMCW ranging system. A linear chirp was applied to the laser light using a silicon photonic IQ modulator fabricated at the University of Southampton, which was driven by a microwave chirp produced by an arbitrary waveform generator (Tektronix AWG70002A). The chirped laser light was amplified by erbium-doped fiber amplifiers (EDFAs) in two stages (a Keopsys CEFA-C-HG-PM followed by an NKT Boostik). The amplified light was then coupled on-chip via single-mode optical fiber V-grooves into two identical demonstration chips, used as a transmitter and receiver respectively. The light emitted by the transmitter FPA was structured using a 32x16 microlens array (PowerPhotonic), which was precisely matched to the receiver array's pixel pattern. This created a structured illumination pattern and minimized any waste of light due to transmit optical power being incident on the gaps between pixels.

To precisely match the fields of view of the transmitter and receiver FPAs, we took advantage of the fact that each receiver grating coupler emits a small amount of LO light due to backreflections from the balanced detectors. An infrared camera was then used to align the patterns of spots produced by the receiver and transmitter FPAs.

The physical aperture of the output lens is 25 mm in diameter. The size of the mode corresponding to a single grating coupler on the receiver focal plane array at this same lens (which can be thought of as the entrance pupil of our system) is slightly elliptical and is 11 mm along the first axis and 16 mm along the second axis. For the long range $75\ \mathrm{m}$ measurements, the output lens was adjusted such that the transmit and receive beams were essentially collimated. For shorter range measurements, the lens position was adjusted to simulate the effects of a smaller aperture.

Thermo-optic switch tree control and calibration. Both the transmit and receive thermo-optic switch trees on our demonstration chip contained integrated photodiodes to monitor the flow of light through the switching trees. To enable digital control and calibration, the monitor photodiodes were directly connected to off-chip TIAs and analog-to-digital converters (Analog Devices AD7091R-8), and the thermo-optic phase shifters were connected to off-chip digital-to-analog converters (Analog Devices AD5391). The switch trees were calibrated one switch at a time by adjusting the control voltage to maximize the optical power in each of the tree outputs. Due to minor thermal cross-talk between the thermo-optic switches, it was necessary to repeat this process for several iterations to converge on an optimal configuration. The average power consumption of a single thermo-optic switch was $40\ \mathrm{mW}$. This leads to an average power consumption of approximately $160\ \mathrm{mW}$ and $120\ \mathrm{mW}$ for the transmitter and receiver switching trees, respectively.

Transmitter chip losses Optical losses in the switching trees are estimated to be $2.1\ \mathrm{dB}$ per switch layer, due in large part to sub-optimal 2×2 couplers. For the transmitter, 4 layers of switches, an additional $1\ \mathrm{dB}$ of waveguide routing loss, and $1.5\ \mathrm{dB}$ of grating coupler loss leads to total linear on-chip losses of $10.9\ \mathrm{dB}$. Under typical transmitter operating conditions, there is an extra $2\ \mathrm{dB}$ loss from two-photon absorption.

These losses can be significantly reduced in a straightforward manner. For example, by employing existing designs for 2×2 couplers with a loss of $0.06\ \mathrm{dB}$ [m3], and thermo-optic phase shifters with a demonstrated loss of $0.23\ \mathrm{dB}$ [m4], the loss could be reduced to $0.35\ \mathrm{dB}$ per switch. Meanwhile, two-photon absorption can be minimized by using larger waveguides or reverse-biased PN junctions[31].

Electronic control and signal processing. A field-programmable gate array (FPGA) with integrated RF ADCs and DACs (Xilinx Zynq UltraScale+ RFSoC) was used for both system control and signal acquisition. Thermo-optic switch control and receiver array multiplexing were coordinated by software running on the system's ARM Cortex-A53 processing core. On the signal processing side, the 8 receiver output signals were first digitized in parallel using 8 integrated ADCs, followed by decimation, application of a Hann window, FFTs, and peak detection on a custom digital signal processing (DSP) pipeline. Final data processing and point cloud reconstruction were performed on a personal computer. To precisely measure the beat frequencies, we performed a least-squares fit of the expected lineshape to each peak in the measured power spectral density. Target distance and velocity were computed using the beat frequencies f_1 and f_2 recorded during the up- and down-chirps respectively. The distance d is given by

$$d = \frac{c(f_1 + f_2)}{4r}, \tag{2}$$

and the velocity v is

$$v = \frac{\lambda_0(f_1 - f_2)}{4}, \tag{3}$$

where r is the chirp ramp rate, c is the speed of light, and λ_0 is the laser wavelength.

Electro-optic characterization A 3.5 GHz oscilloscope (Tektronix DPO7354C) was used for all electro-optic characterization of the receiver array. The amplifier noise floor and shot noise were averaged over a bandwidth of $1 - 3$ MHz, avoiding low frequency $1/f$ noise from the amplifier, as well as the relative intensity noise peak of the laser. The measured shot noise floor was used to determine the exact local oscillator power at each receiver pixel, since the shot noise power spectral density depends only upon the photocurrent. The common mode rejection ratio of the receiver pixels was measured by modulating the amplitude of the local oscillator light at 10 MHz, and comparing the measured electrical output amplitude to the expected amplitude given the local oscillator power and amplifier gain. Based on circuit modelling and independent verification using a test structure on the chip, the total gain of the amplifier chain was 20 kΩ. The total electrical power consumption of the receiver chip was 250 mW. This includes all 8 receiver output channels of signal amplification and multiplexing.

Characterization of measurement accuracy Measurement error in our 3D imaging system can be divided into two categories: systemic errors due to non-idealities in our system, and random fluctuations in the measured beat frequencies due to shot noise, laser relative intensity noise, laser frequency fluctuations, and electronic noise sources. Systemic errors in our system are very tightly controlled. Since the frequency chirps in our system are generated using direct digital synthesis in an AWG, distance accuracy is fundamentally derived from the speed of light, a fixed physical constant, and the timing accuracy of the clocks in the AWG and ADCs, which are controlled to within a few parts per million. The only remaining source of systemic error comes from optical path length differences between pixels, which manifest as static offsets in measured depth. These are due to differences in on-chip optical waveguide lengths, in addition to subtly differing paths taken through the free space optics by light from different pixels. Since these path length differences are static, they can be eliminated using straightforward calibration measurements.

Thus, the key parameter for our system is depth noise, the variation in depth measurements due to stochastic noise in our system. Depth noise was measured by acquiring 40 sequential frames of a static test target for the histograms shown in Fig. 4(b), and 20 sequential frames for the depth precision measurements in Extended Data Fig. 5(b). The mean distance value for each pixel was taken to be the true distance, and depth error was defined as the deviation from the true distance for each pixel. Finally, we defined measurement precision as the standard deviation of the depth error.

Contrast measurement Imaging contrast for the system was measured using 3M Scotchlite 680CR (white) retroreflective sheeting as a target. The edge of the retroreflective sheet was placed across the field of view of the system, allowing for the measurement of pixel-to-pixel crosstalk. The edges were deliberately placed near the middle of the 8×8 pixel receiver blocks to maximize the expected crosstalk. Extended Data Fig. 5(a) illustrates the mean signal strength as a function of distance from the retroreflector edge. This was obtained by averaging the signal strength of 12 pixels in the direction parallel to the retroreflector edge.

The imaging contrast is slightly worse for horizontal edges. This is likely due to the greater numerical aperture of the receiver grating couplers in the vertical direction, leading to additional blurring due to diffraction and lens aberrations in the vertical direction.

Transimpedance amplifier comparison. Fig. 3(d) plots input-referred noise current density against the transimpedance gain-bandwidth product for several state-of-the-art CMOS and BiCMOS optical receiver publications[m5–m9]. A custom design must simultaneously meet requirements for gain, noise, and bandwidth. Generally the gain-bandwidth product will be constant for a target technology and power consumption. In a resistive shunt-feedback configuration, the input-referred noise is typically dominated by the feedback resistor.

$$i_{n,rms} = \sqrt{\frac{4kT}{R_F} \cdot BW_{-3dB}} \tag{4}$$

The gain is approximately equal to the feedback resistance, and the bandwidth is determined by the pole at the input, where C_T is the total capacitance at the TIA input, and A_0 is the open-loop gain of the TIA.

$$BW_{-3dB} = \frac{1 + A_0}{2\pi \cdot R_F \cdot C_T} \tag{5}$$

The negative feedback acts to reduce the input impedance looking into the TIA. Due to our low bandwidth requirement (< 1 GHz), and small diode and parasitic capacitance at the TIA input, we can use a large resistor to get high TIA gain resulting in a reasonable gain-bandwidth product while allowing a low input-referred noise density. Having low-noise electronics improves the system's detection probability, providing longer range for a given optical power.

Methods References

[m1] Stove, A. Linear FMCW radar techniques. *IEE Proceedings F (Radar and Signal Processing)* **139**, 343–350(7) (1992).

[m2] Winkler, V. Range Doppler detection for automotive FMCW radars. In *2007 European Radar Conference*, 166–169 (2007).

[m3] Sheng, Z. *et al.* A compact and low-loss MMI coupler fabricated with CMOS technology. *IEEE Photonics Journal* **4**, 2272–2277 (2012).

[m4] Harris, N. C. *et al.* Efficient, compact and low loss thermo-optic phase shifter in silicon. *Opt. Express* **22**, 10487–10493 (2014).

[m5] Ahmed, M. G. *et al.* 34-GBd linear transimpedance amplifier for 200-Gb/s DP-16-QAM optical coherent receivers. *IEEE Journal of Solid-State Circuits* **54**, 834–844 (2019).

[m6] Shahdoost, S., Medi, A. & Saniei, N. Design of low-noise transimpedance amplifiers with capacitive feedback. *Analog Integrated Circuits and Signal Processing* **86**, 233–240 (2016).

[m7] Mohan, S. S., Hershenson, M. D. M., Boyd, S. P. & Lee, T. H. Bandwidth extension in CMOS with optimized on-chip inductors. *IEEE Journal of Solid-State Circuits* **35**, 346–355 (2000).

[m8] Razavi, B. A 622 Mb/s 4.5 pA//spl radic/Hz CMOS transimpedance amplifier [for optical receiver front-end]. In *2000 IEEE International Solid-State Circuits Conference. Digest of Technical Papers (Cat. No.00CH37056)*, 162–163 (2000).

[m9] Romanova, A. & Barzdenas, V. A review of modern CMOS transimpedance amplifiers for OTDR applications. *Electronics* **8**, 1073–1105 (2019).

Acknowledgements D. J. Thomson acknowledges funding from the Royal Society for his University Research Fellowship. We thank Andy Striker, Andy Watts, Mehrdad Djavid and the Global Foundries team for assistance in device fabrication.

Contributions C.R. and A.Y.P. contributed equally to this work. C.R. conceived, built and tested the free-space portion of the LiDAR system, performed the final LiDAR measurements, and calibrated the optical switching trees. A.Y.P and C.R. performed the electro-optic characterization. A.Y.P. conceived, built and tested the fiber-optic portion of the LiDAR system, wrote the LiDAR system software, tested the coherent receiver array, and contributed to the architecture and layout of the photonic chip. D.J.T. designed and laid out the silicon photonics modulator. R.W. performed design verification including circuit simulations; both created and automated design rule checks for the photonic chip; and contributed to the embedded software control systems. I.O. designed the electronic circuits on the photonic chip, and contributed to their layout. S.A.F developed the signal acquisition and control systems. A.J.C. designed and verified the circuit boards

used to interface with the photonic chip. A.G. contributed to the architecture and performed the layout of the photonic chip. F.M and X.C. contributed to the fabrication and testing of the modulator. R.N. conceived the receiver and switching architecture. R.N. and G.T.R. supervised the project.

Competing Interests In terms of competing interests, all authors with the exception of F.M. and X.C. disclose that they are shareholders of Pointcloud Inc., a start-up company engaged in making laser ranging devices based on coherent receiver arrays.

Correspondence Correspondence should be addressed to R. N. (remus.nicolaescu@point.cloud).

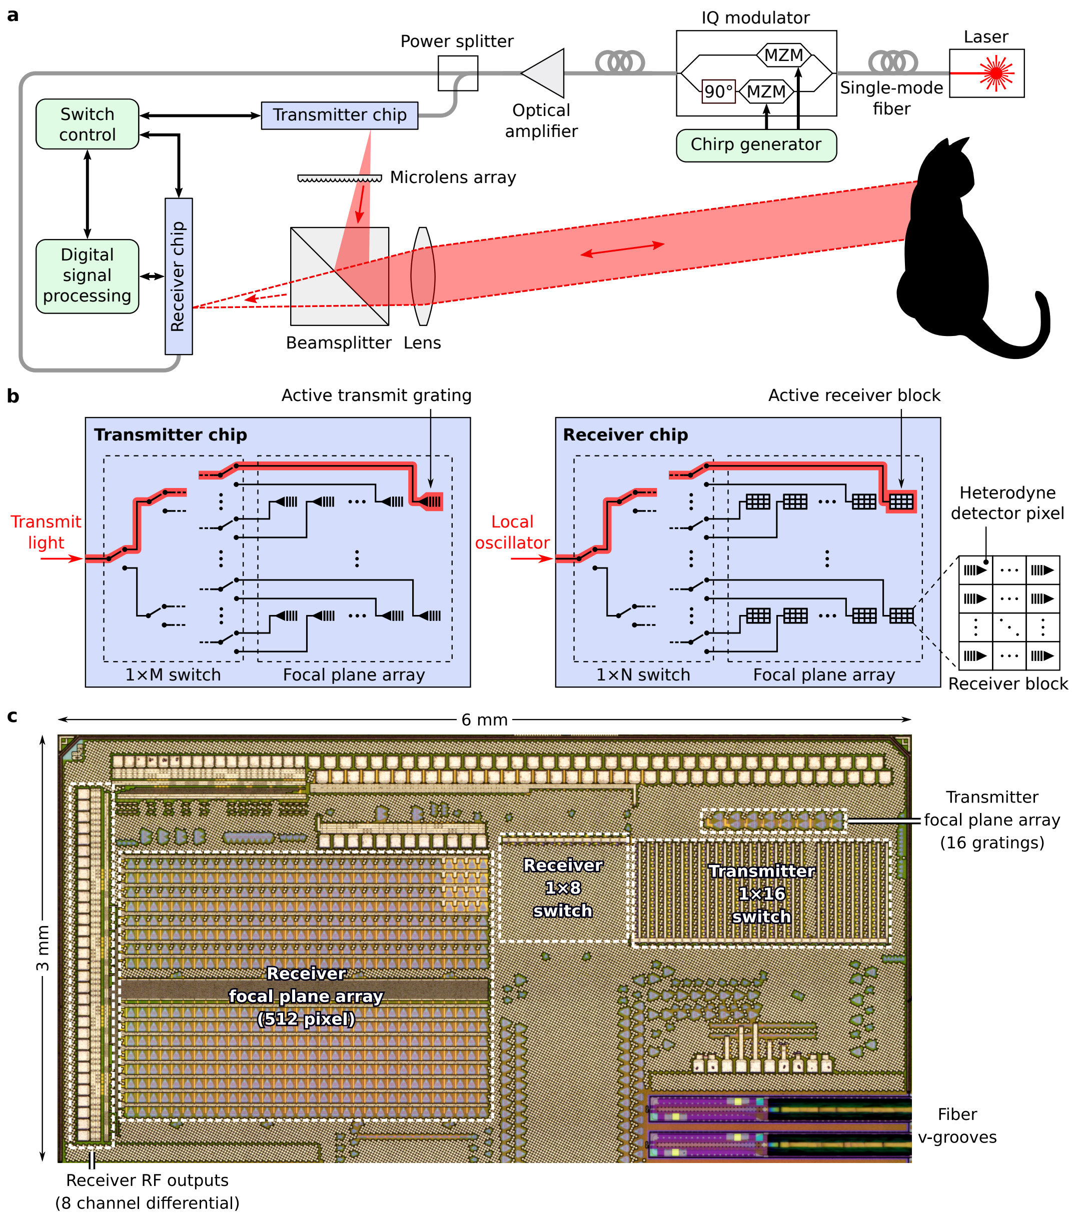

Figure 1: Solid-state 3D imaging architecture. (a) Our architecture consists of two focal plane arrays (FPAs): a transmitter FPA that sequentially illuminates patches of the scene, and a receiver FPA that detects scattered light from the scene. The frequency-modulated continuous-wave (FMCW) scheme is used for ranging. (b) On-chip steering of light is provided by thermo-optic switching trees on both the transmitter and receiver chips. An optional microlens array can be used to shape the illumination pattern to more closely match the receiver array, thereby improving system efficiency. (c) Optical micrograph of our demonstrator chip, showing the switching trees and focal plane arrays for both the transmit and receive functionality.

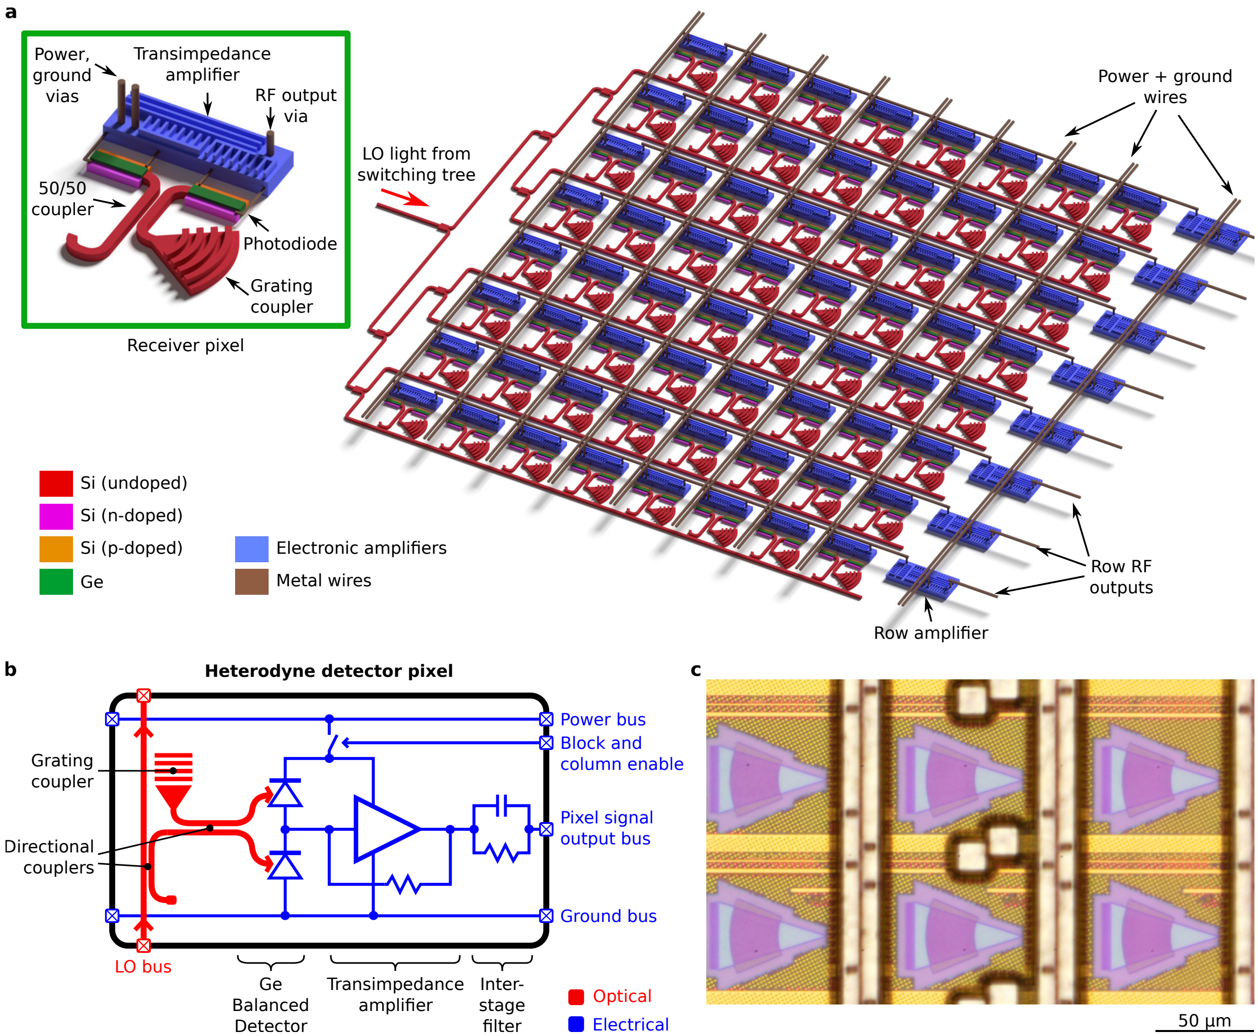

Figure 2: Receiver focal plane array (FPA) design. (a) Schematic of a receiver block in our receiver focal plane array. Within the receiver block, local oscillator (LO) light is distributed to a dense array of heterodyne detector pixels via a network of silicon waveguides. Meanwhile, each pixel collects scattered light from the scene using a grating coupler, which is combined with LO light on a balanced detector to produce a detectable photocurrent. The photocurrent is amplified in two stages: first by a transimpedance amplifier (TIA) within the pixel, and again by an amplifier at the end of each row. For clarity, we have omitted control wires from the diagram. (b) Electrical schematic of the heterodyne detector pixel. (c) Optical micrograph of a small subset of the receiver focal plane array.

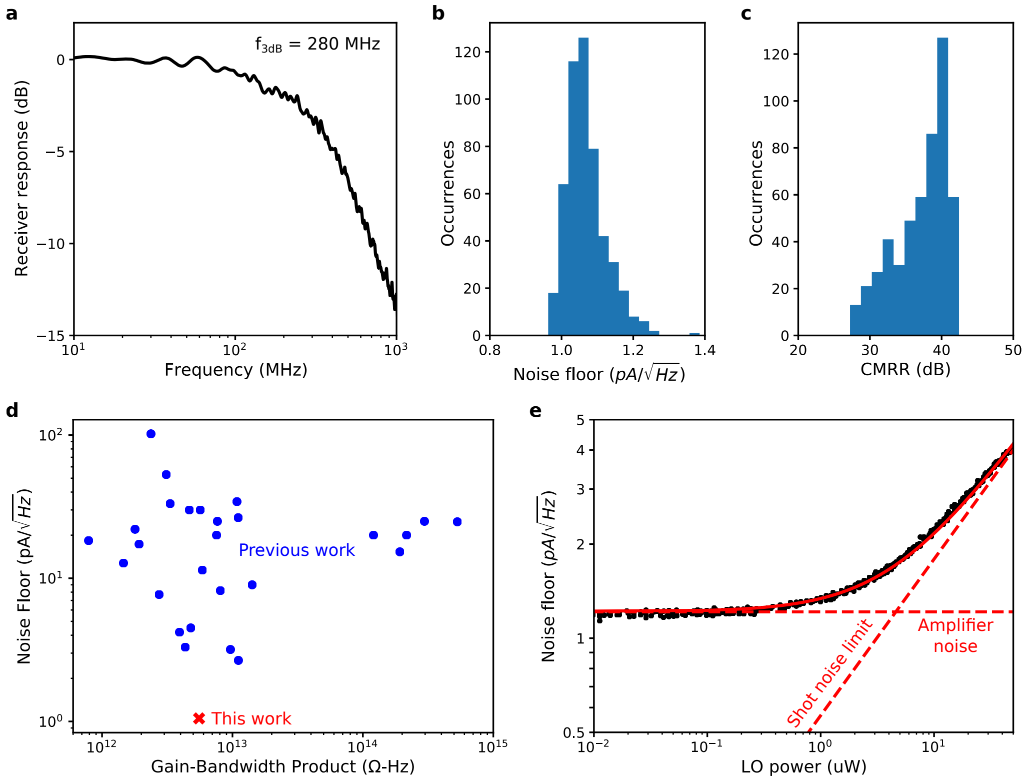

Figure 3: Receiver electro-optic performance. (a) Measured frequency response of the receiver readout chain for an optical signal supplied to a single pixel, showing a cutoff frequency of $280\ \mathrm{MHz}$. (b, c) Histograms of input-referred amplifier noise and common-mode rejection ratio (CMRR) respectively throughout the full array, showing tight distributions for both parameters. (d) Largely due to tight integration between our photodiodes and TIAs, we have achieved a high gain-bandwidth product with significantly improved noise performance compared to previous designs. (e) Input-referred noise as a function of optical local oscillator (LO) power for a single pixel, demonstrating shot-noise limited detection using $< 10\ \mu\mathrm{W}$ of LO power.

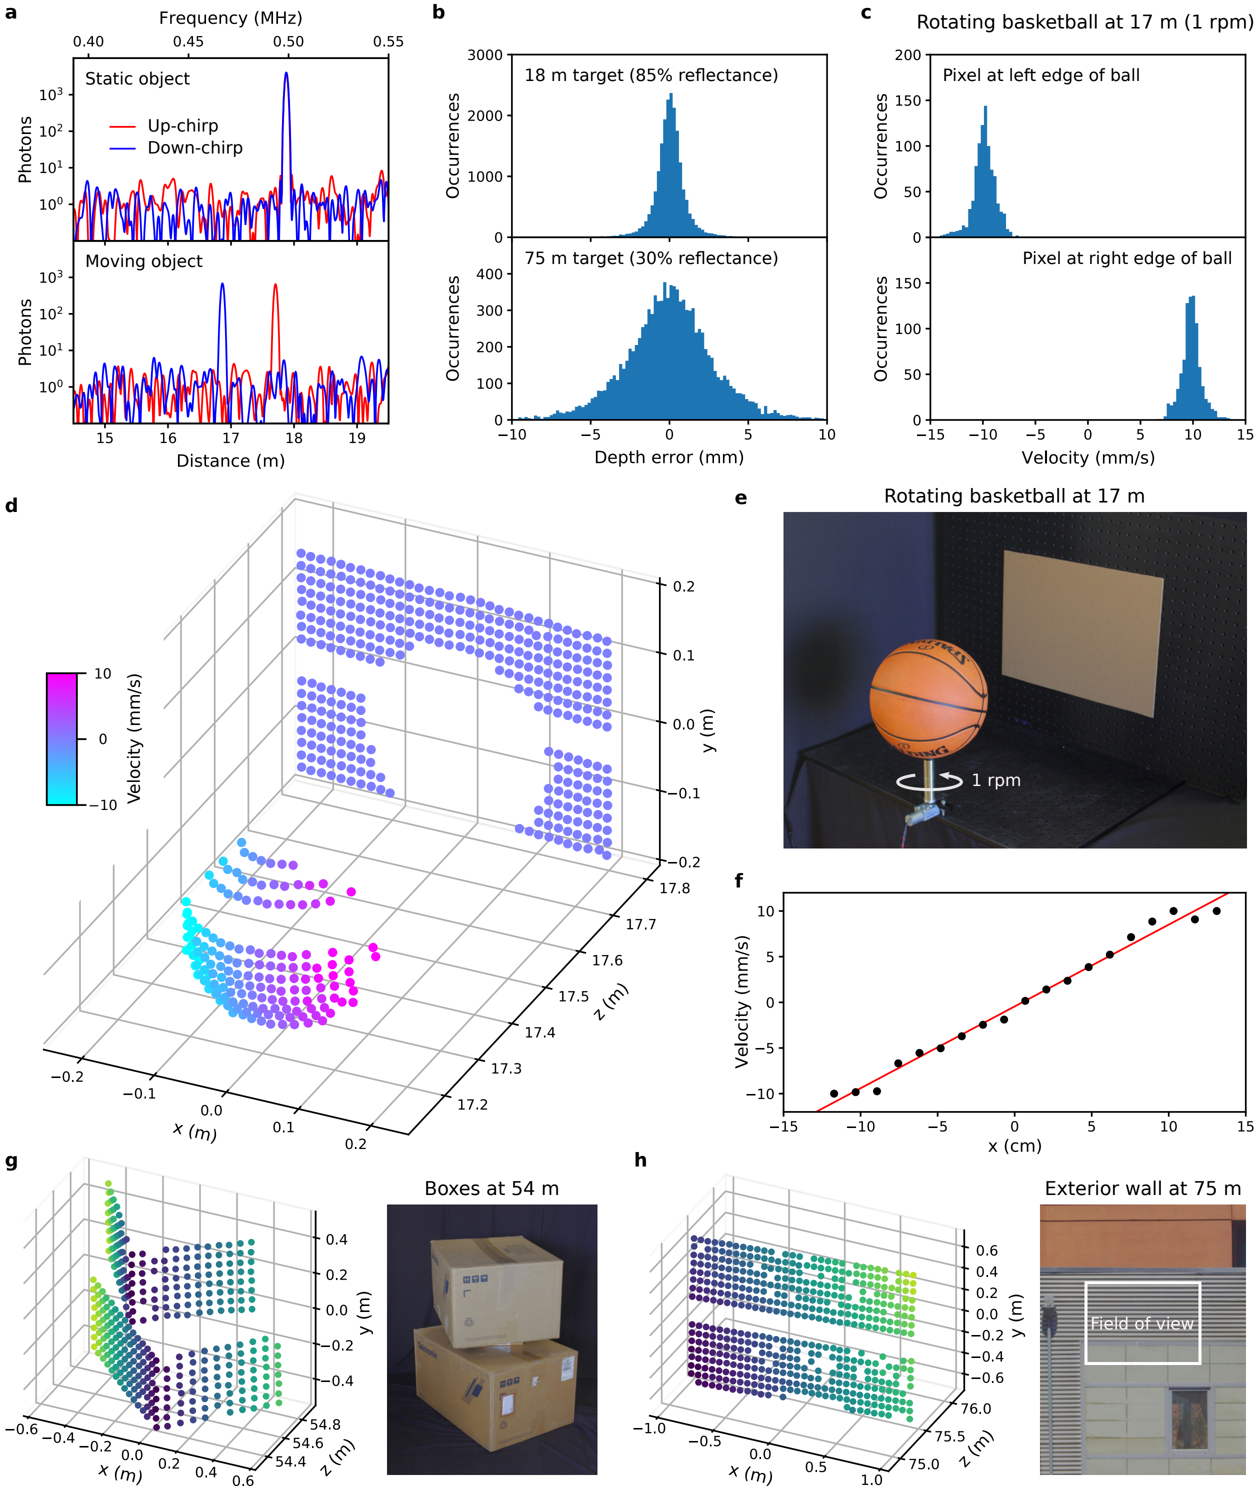

Figure 4: 3D imaging system characterization. (a) Representative signals from a receiver pixel, showing Doppler splitting between the up- and down-chirps for the moving target. (b) Depth noise for targets at $18\ \mathrm{m}$ and $75\ \mathrm{m}$, with standard deviations of $1.8\ \mathrm{mm}$ and $3.1\ \mathrm{mm}$ respectively. (c) Velocity histograms for a basketball rotating at $1\ \mathrm{rpm}$, exhibiting a standard deviation of $1.0\ \mathrm{mm/s}$. (d) Velocity annotated point cloud of a basketball at $17\ \mathrm{m}$ rotating about its vertical axis at $1\ \mathrm{rpm}$. (e) Photograph of the basketball setup. (f) Horizontal linecut of velocity across the middle of the basketball. (g, h) Point clouds of (g) stacked cardboard boxes at $54\ \mathrm{m}$, and (h) an exterior wall at $75\ \mathrm{m}$. Distance to the target is indicated by colour in (e) and (f). The missing band of points in the middle of the point clouds is due to a narrow gap in the receiver array for electrical and optical routing.

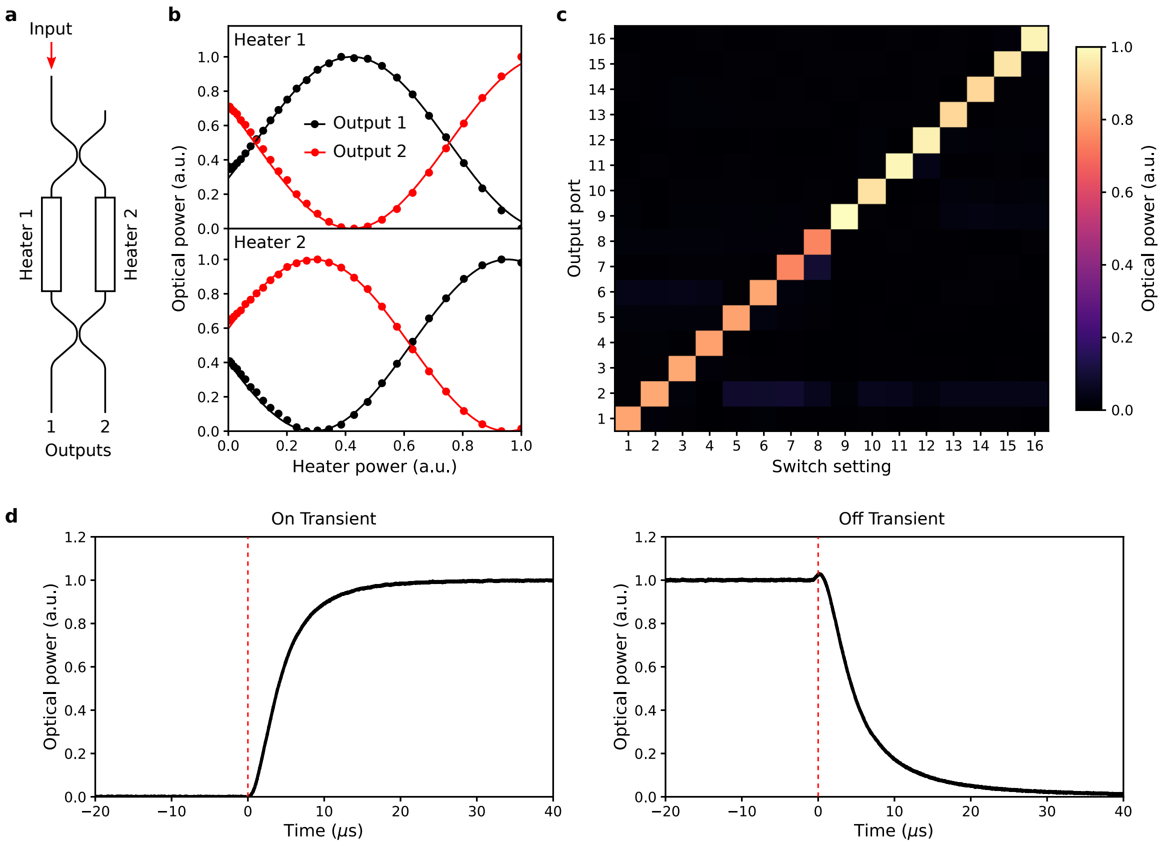

Extended Data Figure 1: Thermo-optic switching tree demonstration. (a) The thermo-optic switches consist of a Mach-Zehnder interferometer with an electrical heater on each arm. (b) Tuning curve for a single thermo-optic switch, showing optical power in the two outputs as a function of applied heater power. The use of two heaters allows the average electrical power consumption per switch to be halved. (c) Output power distribution of the 1×16 transmitter switch tree for all switch settings, demonstrating clean switching. Output power was monitored using a set of monitor photodiodes at the output of the switch tree. (d) On and off transients for a representative thermo-optic switch, demonstrating $90\% - 10\%$ switching times of 9.1 μs and 12.1 μs respectively. Due to minor thermal crosstalk between switches, the switching transients are not perfect decaying exponentials.

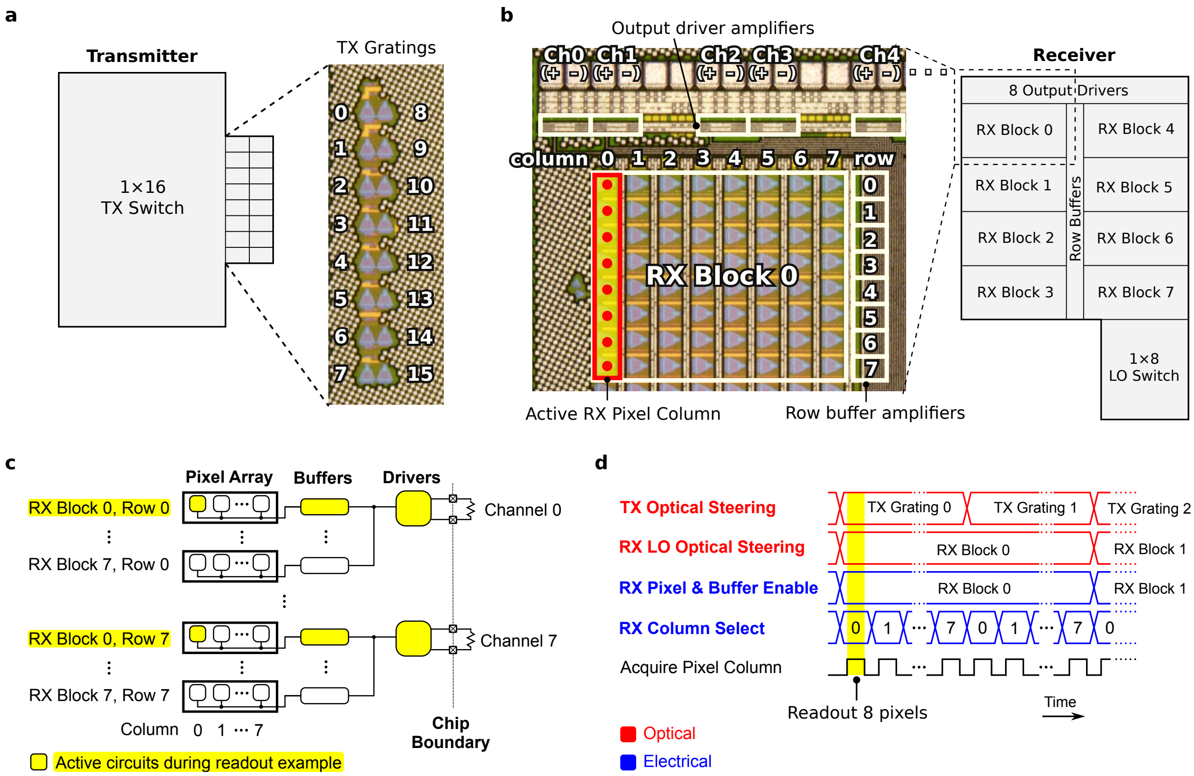

Extended Data Figure 2: Transmitter (TX) and receiver (RX) synchronization and readout architecture. (a) The TX steers light through a 4 level tree of 1×2 switches to feed the FPA of 16 output grating couplers. Each leaf contains a fractional tap and monitor photodiode enabling electronic calibration of the tree. (b) The RX array is divided into 8 blocks of 64 pixels. Imaging an 8-pixel column requires both steering the LO light to the block and enabling the associated electronics (pixel column and row buffer amplifiers). Signals from the active pixel column are driven by 8 output amplifiers for parallel readout. (c) Several levels of multiplexing are used to map 512 pixels down to 8 output channels. An active RX block has one active pixel per row, with the other disabled pixels within the row presenting high output impedance (no drive strength). The row buffers are similarly passively multiplexed between the blocks. The 8 drivers are always activated during readout. (d) Timing diagram showing synchronization between the optical switching trees (TX and RX) and the electrical readout circuitry.

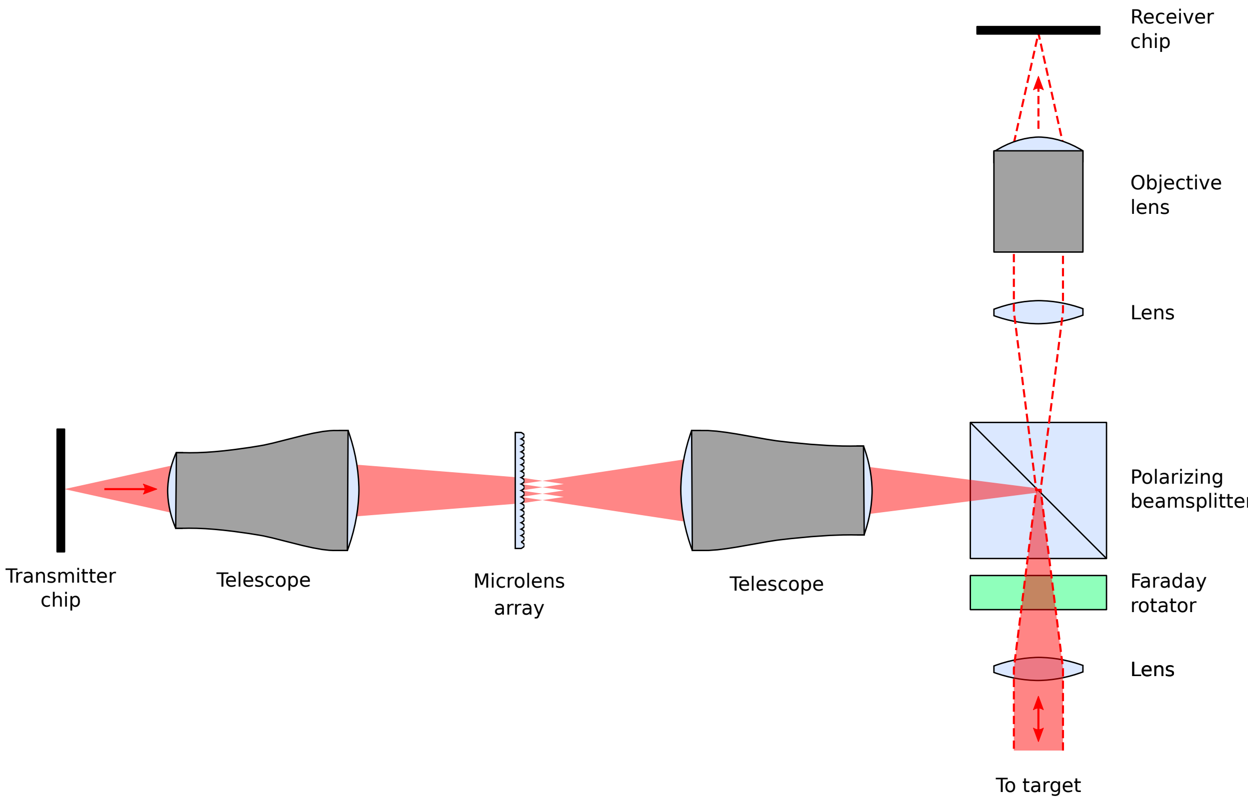

Extended Data Figure 3: Free-space optics schematic of the demonstration system. Much of the complexity in the optical system is to match the receiver and transmitter focal plane arrays, which can be corrected in the future by adjusting the on-chip layouts. For inexpensive consumer versions of the system, the Faraday rotator and polarizing beamsplitter could be replaced by a 50-50 beamsplitter, at the cost of a $4\times$ reduction in signal strength. Although it is possible to implement this experiment using a single chip for both transmit and receive functions, we have used two identical chips acting as the transmitter and receiver respectively to simplify the experimental setup.

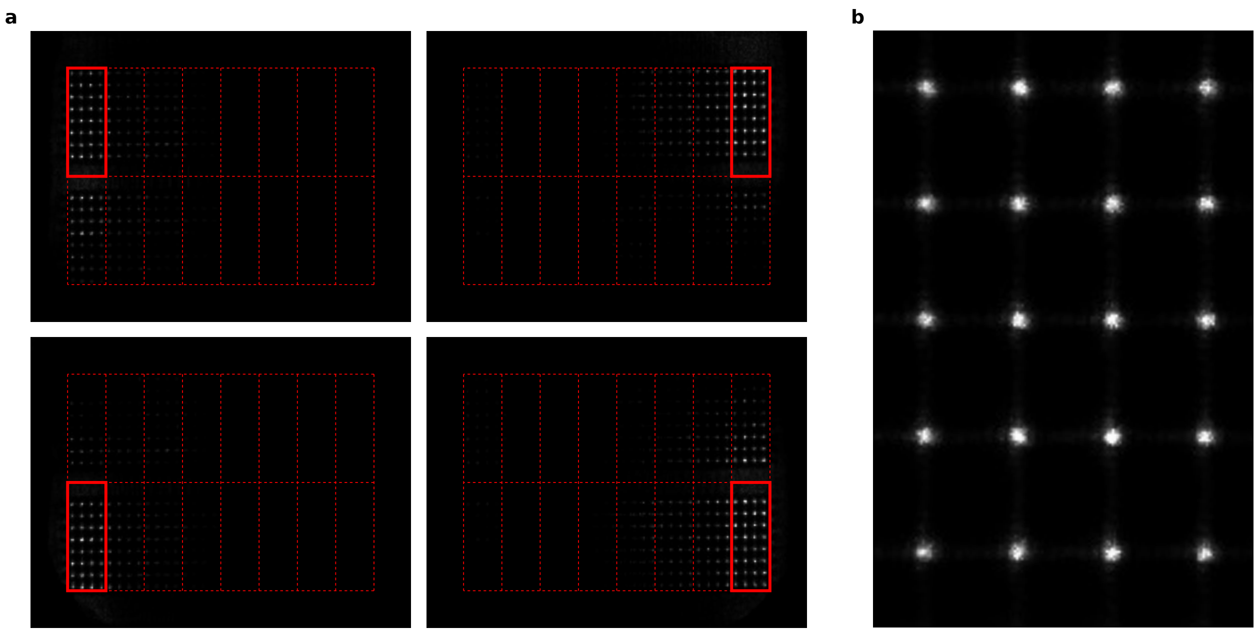

Extended Data Figure 4: Far-field infrared camera images of transmitter steering. (a) Images of several representative steering positions. The receiver fields of view corresponding to the 16 steering positions are indicated by the dashed lines, with the currently active block indicated by a solid outline. The light from the active transmit grating is first slightly defocused to completely illuminate the active block, and then structured by the microlens array. Due to this defocusing, a small fraction of the transmitted light falls outside of the active block. (b) A zoomed-in image showing the structured illumination pattern produced by the microlens array. The locations of the bright spots coincide with the receiver pixel grating couplers.

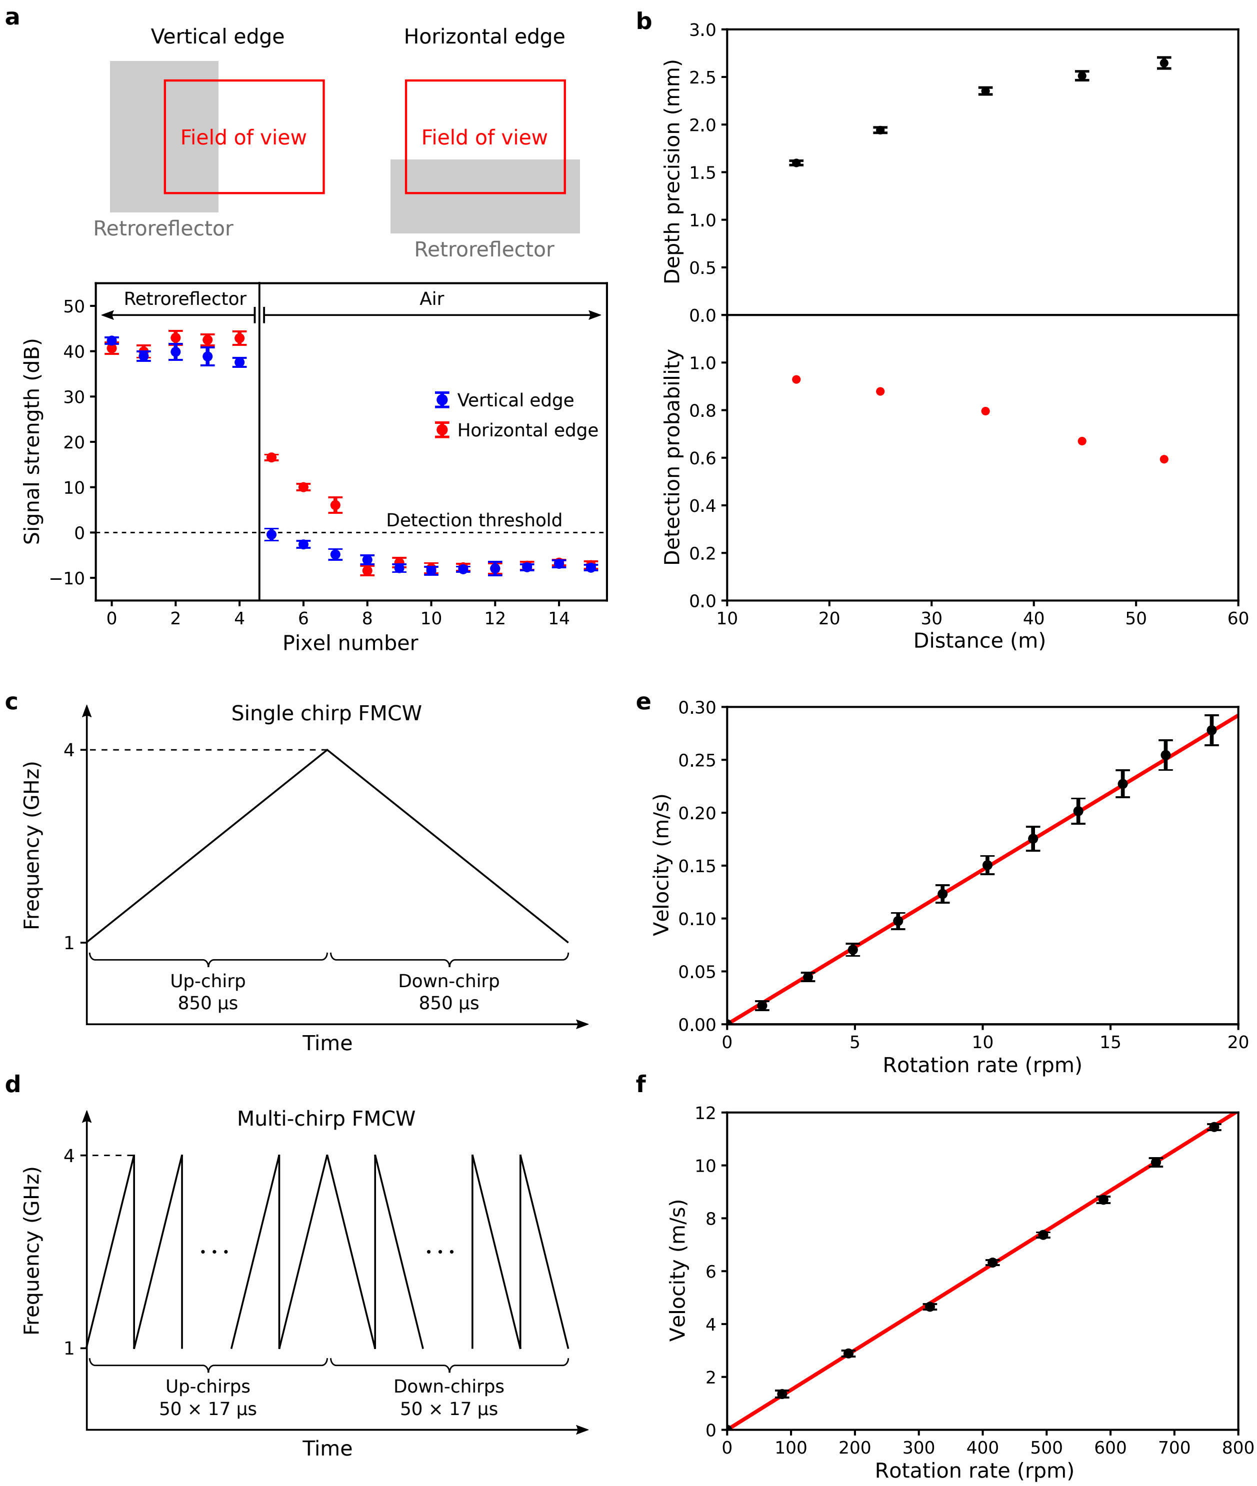

Extended Data Figure 5: Additional characterization of system performance. (a) Imaging contrast measured using retroreflective sheeting. Our system achieves > 25 dB contrast for a 1 pixel displacement, > 50 dB contrast for a 4 pixel displacement, and reaches the system noise floor thereafter, illustrating the excellent pixel-to-pixel isolation in our system. Here, the error bars represent the standard error. (b) Depth precision and detection probability as a function of distance for a 44% reflectance target. The error bars on the depth precision represent the 95% confidence intervals. (c, d) Single-chirp (c) and multi-chirp (d) FMCW waveforms used for measuring slow and fast objects respectively. (e, f) Measured velocity as a function of rotation rate for a 30 cm diameter styrofoam cylinder at a distance of 17 m using (e) single-chirp and (f) multi-chirp waveforms, with error bars indicating the standard deviation.