A 256-Element CMOS Imaging Receiver for Free-Space Optical Communication

Brian S. Leibowitz, *Member, IEEE*, Bernhard E. Boser, *Fellow, IEEE*, and Kristofer S. J. Pister

***Abstract*—This paper describes a CMOS imaging receiver for free-space optical (FSO) communication. The die contains 256 optical receive channels with −47 dBm optical sensitivity and 30 dB optical dynamic range at 500 kb/s/channel while consuming 67 mW. Received signals are amplified by digitally self-calibrated open-loop amplifiers and digitized before clock and data recovery. The sampled data also provide inputs for digital automatic gain and offset control loops closed around the analog amplifier chain to compensate for signal variations due to atmospheric turbulence and daylight interference. Gain control logic can adapt to incident signals over the 30 dB dynamic range within 28 bit periods. Low-power logic design and analog circuit techniques are used to minimize digital crosstalk to single-ended photodetectors referenced to a bulk substrate. Local arbitration circuitry at each channel forms an intrachip data passing network to multiplex received data words from the 16 × 16 array onto a common off-chip bus. The 1.6 M transistor mixed-signal die fabricated in a 0.25 μm CMOS process measures 6.5 × 6.5 mm^2. Reception at 500 kb/s through a 1.5 km atmospheric channel is demonstrated with 3 mW optical transmit power during nighttime and daylight hours.**

***Index Terms*—Free-space optical communication, gain control, integrated optoelectronics, optical arrays, optical receivers, self calibration.**

I. Introduction

FREE-SPACE optical (FSO) is an attractive alternative to RF for long-range communication between small low-power devices because the shorter radiation wavelengths utilized allow for more directional transceivers. For example, visible light can be collimated to approximately 1 mrad divergence by a 1 mm lens, which would require an antenna several hundred meters in diameter for typical RF wavelengths of several centimeters. Such a highly collimated transmission has an antenna gain greater than 60 dB, which directly reduces the transmit power requirement for a long communication link. An additional benefit of highly directional transmissions is that nonaligned transceivers do not interfere with one another. This allows for multiple-access through spatial division (SDMA), without the need for code-, time-, or frequency-division architectures, at an implicit cost that the transceivers must be precisely aligned. FSO communication also has incidental benefits, such as lack of government regulation and use of baseband signaling with generally relaxed linearity requirements compared to most RF modulation schemes. These properties of FSO communication make it attractive in applications such as distributed sensor networks, "last mile" connectivity in rural areas, and custom high-speed point-point links such as the one described in this paper.

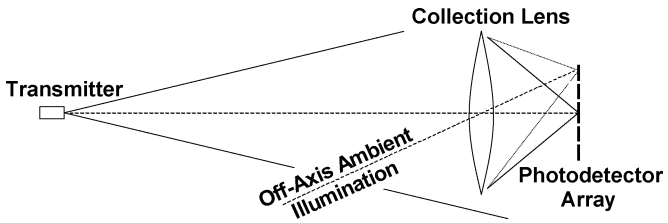

Fig. 1. Imaging receiver uses a lens to discriminate received light based on angle of incidence, separating the desired signal from most ambient illumination.

The most notable drawback to FSO communication is that direct line of sight is required for long links. Furthermore, in daytime links, ambient daylight acts as a strong low-frequency interfering signal for an FSO receiver. Although this interference signal may be rejected by an electronic high-pass filter, it generates wideband shot noise at the photodetector that fundamentally limits noise performance. In addition to narrowband optical filtering, ambient interference can be reduced with an imaging receiver consisting of a lens and an array of photodetectors as shown in Fig. 1 [1], [2]. Ambient light received from the full field of view is distributed across the detector array based on the angle of incidence. Each individual detector only receives a fraction of the collected ambient light, while the collected signal power is ideally focused onto a single photodiode. High array resolution is therefore desirable to improve the optical SNR, but interconnect challenges make this prohibitive for custom photodetector arrays with external receiver electronics. Integration of the detector array and associated receiver electronics into a single CMOS die provides a way to achieve high array resolution and high data rates simultaneously.

This paper describes a 256-element CMOS imaging receiver, with a photodetector and receiver electronics integrated at every element, forming a fully parallel array of optical receivers. This receiver is designed for communication between small unmanned aerial vehicles (UAVs) performing visual ground surveillance [3]. In this application, size and power are tightly constrained, and line of sight is implicitly available, making FSO an attractive solution. Link ranges of 100 m–5 km are desired at an altitude greater than 100 m, requiring 34 dB optical dynamic range at the receiver, assuming fixed transmission power. Complete discussion of the link design involves numerous tradeoffs between transmitter,

Manuscript received December 8, 2004; revised February 23, 2005.

B. S. Leibowitz is with Rambus, Inc., Los Altos, CA 94022 USA (e-mail: leibowitz@ieee.org).

B. E. Boser and K. S. J. Pister are with the Berkeley Sensor and Actuator Center University of California, Berkeley, CA 94720 USA (e-mail: boser@eecs.berkeley.edu; pister@eecs.berkeley.edu).

Digital Object Identifier 10.1109/JSSC.2005.848148

receiver, and beam-alignment control systems across optical, mechanical, and electronic domains and is beyond the scope of this paper. Codesign of these systems, subject to constraints such as laser eye safety concerns, spectral responsivity of available silicon photodiodes, and optical form factors at both ends of the link, lead to a designed transmission wavelength of 660 nm with less than 5 mW average power (similar to the beam produced by a handheld laser pointer) and a minimum received optical signal power of −46 dBm (25 nW) with up to 250 nW of interference from ambient illumination. In addition, for transceivers to acquire initial link alignment in a dynamic environment, the receiver must adapt to newly incident signals over the full dynamic range and begin receiving data within 100 μS, or within 50 bit periods at 500 kb/s.

For the link ranges in this application, atmospheric turbulence leads to time-varying channel attenuation that is tracked and compensated for by the receiver. Atmospheric turbulence is highly variable, but, based on an empirical measurement of turbulence strength under conditions similar to those of the intended application, we predict that turbulence in our application will lead to time-varying channel attenuation on the order of 10 dB with millisecond time constants [4], [5]. Because turbulence is much stronger near ground level, the achievable range for ground–ground links will be lower than that of air–air links.

The following section describes the architecture and circuit implementation of an optical receiver designed to meet these dynamic range, sensitivity, and turbulent scintillation tolerance requirements, with low enough power consumption and complexity to allow for array integration. Mixed-signal issues pertaining to array integration and circuit solutions are discussed. Measured performance of the fabricated 256-element imaging receiver confirms the designed sensitivity and dynamic range. Measured results for a 1.5 km ground–ground FSO link are presented, demonstrating that long-range, low-power FSO links are realizable with such an imaging receiver, despite the effects of atmospheric turbulence and interference from ambient daylight and the constraints of array integration.

II. Receiver Design and Performance

A. Single-Channel Receiver Architecture

Each of the 256 parallel optical receive channels are based on the architecture shown in Fig. 2. All analog signals are fully differential, but are shown single-ended for clarity. The detected photocurrent is sensed by a series impedance, amplified by an LNA and a programmable gain amplifier (PGA), and then oversampled by a coarse flash ADC with four comparators. Results of the ADC conversion are monitored by a finite state machine (FSM) that closes digital automatic gain and offset control (AGC/AOC) loops around the signal chain to properly condition the signal swing at the PGA output for data slicing. The oversampled binary data stream is then processed by an FSM that performs digital clock and data recovery (CDR) and data frame detection to produce received data words. Framed data words are finally passed through neighboring receivers in the array toward a shared off-chip bus driver.

Optical signals are on/off keyed (OOK) and Manchester coded to reduce low-frequency signal content that would be masked by ambient illumination and to allow for a simple digital clock-recovery implementation similar to that presented in [6]. In addition, Manchester coding ensures that both "on" and "off" signal levels can be observed within any 1.5-symbol period, allowing gain and offset control updates to be applied periodically at fixed intervals, independent of both signal data and clock recovery status.

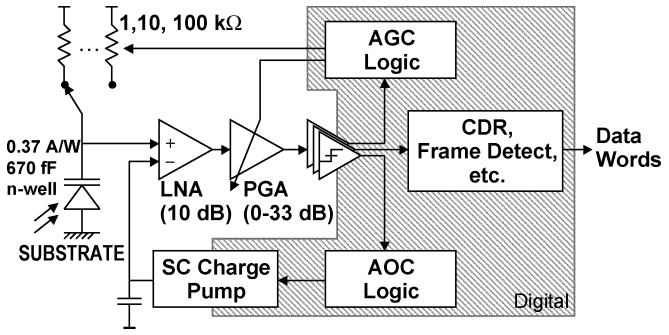

Fig. 2. Single-channel optical receiver architecture.

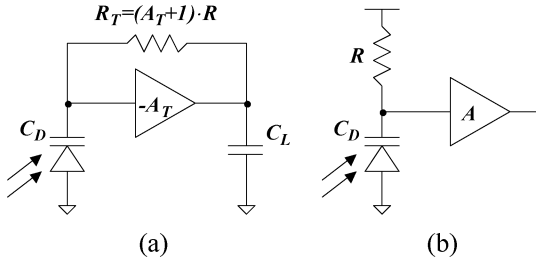

Fig. 3. (a) Transimpedance amplifier topology. (b) Series-impedance amplifier topology.

Key components of this receiver architecture are described in the following sections.

B. Front-End Detection

Optical signals are detected by a reverse-biased n-well junction diode that generates a photocurrent proportional to incident optical power. The optical responsivity of this junction in the CMOS process utilized is 0.36 A/W at 660 nm. The −46 dBm minimum detectable optical signal and 34 dB dynamic range requirements specified above translate to a minimum signal current of 9 nA, a dynamic range of 68 dB (because optical power is converted linearly to signal current), and a superposed ambient interference current of up to 100 nA.

The transimpedance amplifier, shown in Fig. 3(a), is widely used in optical receivers for several reasons, but the primary advantage over the series impedance topology shown in Fig. 3(b) is that it can use approximately $A_T + 1$ times higher resistance for an equivalent bandwidth and input capacitance, thereby reducing the thermal noise current contribution from the resistor, at the expense of a smaller dynamic range [7]. However, in a sufficiently power-constrained application, the noise contribution of the amplifier dominates that of the feedback resistor, limiting the achievable noise benefit of the transimpedance topology, as discussed below. Furthermore, the dynamic range of the series-impedance topology can be extended in a straightforward manner through use of a nonlinear or programmable load impedance.

Consider the noise performance of the series-impedance topology with a wideband and noiseless amplifier with sense resistance R and input capacitance C (typically dominated by the photodiode capacitance C_D) which determine the signal bandwidth $B = (2\pi \mathrm{RC})^{-1}$. In this ideal scenario, the input referred noise current is

$$\overline{i_r^2} = \frac{4kT}{R}\frac{\pi}{2}B = 4\pi^2 kT\, B^2 C = 2\pi kT\, B|Y| \qquad (1)$$

where k is Boltzmann's constant, T is the absolute temperature, and $|Y| = 2\pi BC$ is an admittance magnitude that is generally set by application requirements and photodetector process technology.

In a power-constrained scenario, the voltage amplifier degrades both noise performance and signal bandwidth, requiring the RC input and amplifier pole frequencies to both be greater than B. In the extreme case when the amplifier is the dominant noise source, the amplifier pole will be set to $f_a \approx B$ and the input node will be set relatively wideband. In this case, the input-referred thermal noise current contribution from the MOS amplifier is

$$\overline{i_a^2} = \frac{8}{3}\frac{kT}{g_m^* R^2}\frac{\pi}{2}B = \frac{16}{3}\frac{kT}{g_m^*}\pi^3 B^3 C^2 = \frac{8}{3}\frac{kT}{g_m^*}\pi^2 B^2 C|Y| \qquad (2)$$

where g_m^* is the noise equivalent transconductance of the voltage amplifier, typically similar to the transconductance of the input transistor. Flicker noise can be made negligible in most cases where a large photodiode capacitance allows use of relatively large gate-area devices without significant degradation in input node capacitance. The sense resistor and amplifier noise contributions (1) and (2) are equal when $g_m^* = (2/3)|Y|$. Noise performance will be limited by resistor thermal noise or by amplifier noise for much higher or lower values of g_m^*, respectively.

For the transimpedance topology with an ideal wideband and noiseless amplifier, the transimpedance R_T, input capacitance C, and amplifier gain A_T determine the noise and signal bandwidth $B = (A_T + 1)/(2\pi R_T C)$. In this scenario, the input referred noise current is

$$\overline{i_{rT}^2} = \frac{4kT}{R_T}\frac{\pi}{2}B = \frac{4\pi^2 kT\, B^2 C}{A_T + 1} = 2\pi\frac{kT\, B|Y|}{A_T + 1}. \qquad (3)$$

Slightly better noise performance can be achieved with a finite-bandwidth amplifier by designing for maximally flat frequency response (taking advantage of complex pole splitting to introduce a second pole into the noise transfer function and increase the -3 dB signal bandwidth), but this requires the voltage amplifier pole frequency to be accurately set to $2(A_T + 1)/(2\pi R_T C)$ and will not be considered here.

In a power-constrained transimpedance topology, voltage amplifier noise degrades the overall noise performance, as in the case of the series-impedance topology, but finite amplifier bandwidth also affects the stability of the transimpedance loop. To guarantee stability with greater than 60° phase margin, without designing for a maximally flat response as discussed above, the transimpedance amplifier is generally designed with $R_T \approx A_T/(2\pi BC)$ and voltage amplifier pole frequency $f_a > 2B$.

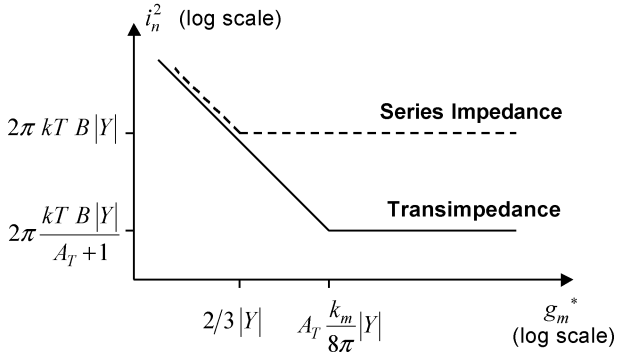

Fig. 4. Transimpedance and series-impedance input-referred noise power versus g_m^*.

In this case, the resistor noise contribution is approximately unchanged from (3) and the input-referred MOS amplifier noise current spectral density (the output noise voltage spectral density divided by the transimpedance) is

$$\frac{\overline{i_{aT}^2}}{\Delta f} \approx \frac{8}{3}\frac{kT}{g_m^* R_T^2}\left(\frac{1 + sR_T C_D}{1 + s\frac{R_T C_D}{A_T}}\right)^2 \left(\frac{1}{1 + \frac{s}{2\pi f_a}}\right)^2 \qquad (4)$$

under the assumptions $A_T \gg 1$ and $f_a \gg B/A_T$. The resulting integrated mean square noise current can be approximated by

$$\overline{i_{aT}^2} \approx k_m \pi^2 \frac{kT}{g_m^*} B^3 C_D^2 = k_m \frac{\pi}{2}\frac{kT}{g_m^*} B^2 C|Y| \qquad (5)$$

where k_m is a dimensionless constant that is primarily dependent on the loop phase margin (on the excess bandwidth of the voltage amplifier). Numerical integration of (4) yields $k_m \approx 20$ for a transimpedance amplifier with 60° phase margin and $A_T = 10$.

Combining (3) and (5), the input-referred noise current for the transimpedance topology is

$$\overline{i_{nT}^2} \approx \overline{i_{rT}^2} + \overline{i_{aT}^2} \approx \frac{2\pi}{A_T} kT\, B|Y|\left(1 + A_T \frac{k_m}{8\pi}\frac{|Y|}{g_m^*}\right). \qquad (6)$$

The term outside the parentheses represents the feedback resistor noise contribution, and the term inside the parentheses is the noise factor added by the voltage amplifier with finite power consumption.

The input-referred noise of the series-impedance and transimpedance topologies are compared in Fig. 4, using g_m^* as a proxy for power consumption. For large g_m^*, the transimpedance topology achieves a factor $A_T + 1$ lower input-referred noise power. However, in the most power-constrained case with $g_m^* < |Y|$, the two topologies have similar input-referred noise powers regardless of A_T.

In the receiver described here, power constraints limit g_m to approximately 100 μS, with a desired front-end bandwidth of 2.4 MHz. The photodiode used is a 125×125 μm^2 n-well junction, which has a parasitic capacitance of 0.67 pF at a 2.5 V reverse bias. This constrains $g_m^* < 10|Y|$, since g_m^* is generally smaller than g_m. Under these constraints, the transimpedance topology offers at most a 10 dB noise advantage over the series-impedance topology.

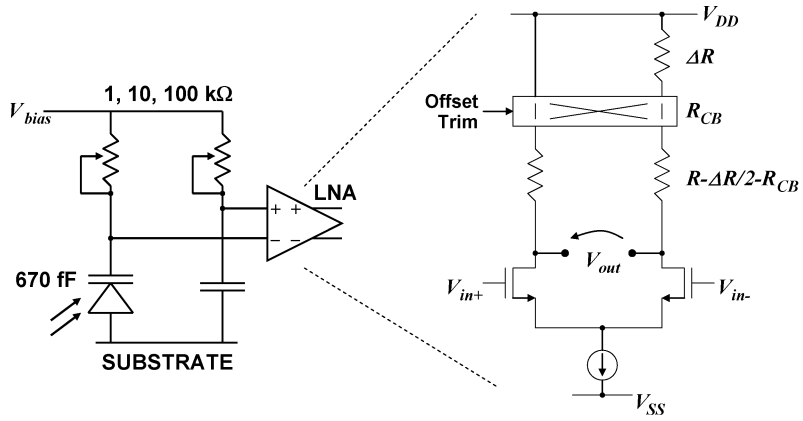

Fig. 5. Open-loop differential LNA topology with only one of three digital offset trim crossbar switches shown.

In this application, this limited noise benefit of the transimpedance topology is sacrificed to take advantage of other benefits of the series-impedance topology, primarily the wider dynamic range of the series-impedance topology and the ability to further extend the dynamic range with a nonlinear or variable sense impedance without stability concerns. In this case, the receiver dynamic range is extended by 40 dB with a digitally programmable sense resistance of 1, 10, or 100 kΩ, as shown in Fig. 5, and controlled as described in Section II-C.

As shown in Fig. 5, the voltage signal developed across the sense impedance is amplified by a differential LNA with a replica *RC* network connected to the negative input. This replica network is described in greater detail in Section II-F. The LNA is designed for fixed 10 dB gain and 67-μV rms input-referred noise, including flicker noise. The LNA consumes 24 μA from a 2.5 V supply, which was chosen to accommodate a common-mode input level at the maximum photodiode bias voltage $V_{\rm bias} = 2.5$ V. Three pMOS crossbar switches are used to adjust binary weighted trim resistors in the differential load to coarsely adjust the fixed amplifier offset voltage using a built-in self-calibration routine described in Section II-D below. Maximum sensitivity is achieved with 100 kΩ front-end sense impedance, where the LNA input noise voltage is equivalent to an rms input noise current of 0.67 nA. The front-end pole at 2.4 MHz and a combined LNA/PGA bandwidth of 4.4 MHz lead to a front-end signal bandwidth of approximately 1.6 MHz. Over this bandwidth, thermal noise from the two 100 kΩ input resistors leads to an rms noise current of 0.89 nA, and the maximum dc photocurrent of 100 nA from ambient interference leads to an rms shot noise current of 0.34 nA. The designed input-referred noise current from these three sources is 1.16 nA at a density of 0.93 pA/$\sqrt{\rm Hz}$. With photodiode responsivity 0.36 A/W, this corresponds to an optical noise equivalent power (NEP) of −55 dBm.

Fig. 6 shows an eye diagram measured at the PGA output of a single receiver element with a 19-nW (−47 dBm) average optical power test signal, which is the minimum level that can be reliably received. This generates an expected 7 nA average (14 $\rm nA_{pp}$) photocurrent that is sensed by the maximum front-end impedance of 100 kΩ, resulting in 15 dB SNR based on the designed noise performance discussed above. The test signal was applied directly to a single receiver in the array via an optical fiber positioned directly above the desired photodiode. Test signal power was determined by measuring the total fiber exit beam power. In this measurement, ambient illumination was adjusted to generate a 100 nA dc front-end signal current in addition to the test signal to include the effect of ambient-induced shot noise. The dc ambient signal itself is subtracted by an automatic offset control loop described below.

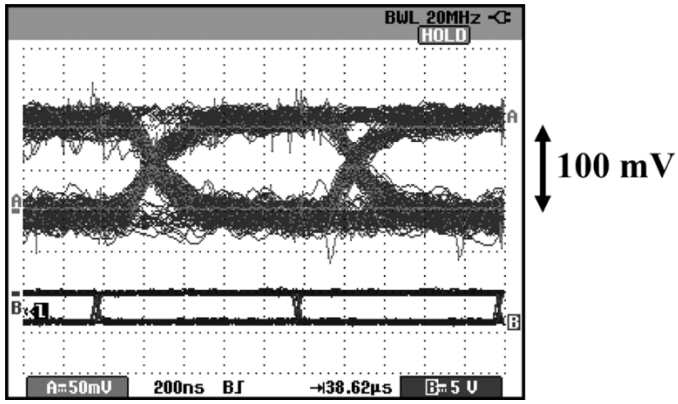

Fig. 6. Received eye diagram (top trace) for 19-nW average, 38-$\rm nW_{pp}$ optical signal power at 500 kb/s with pseudorandom Manchester coded test data (bottom trace).

C. Automatic Gain and Offset Control

Automatic gain and offset control loops are used to accommodate wide dynamic range inputs and to correct signal variations caused by time-varying channel attenuation due to atmospheric turbulence and varying ambient illumination. The AGC loop uses an adaptive two-phase search algorithm based on coarse gain control implemented through selectable front-end sense impedances of 1, 10, and 100 kΩ and fine gain control implemented through a 0–33 dB PGA with 3 dB step size, which is described in Section II-D. The PGA control range is made wider than 20 dB to provide hysteresis at the coarse gain control boundaries. The AOC loop is based on an up/down switched-capacitor charge pump derived from the topology in [8]. The charge pump output drives a differential transconductor to subtract current from a summing node at the LNA output. This transconductor uses one-sixth as much bias current as the LNA and is resistively degenerated to provide 20 dB lower gain

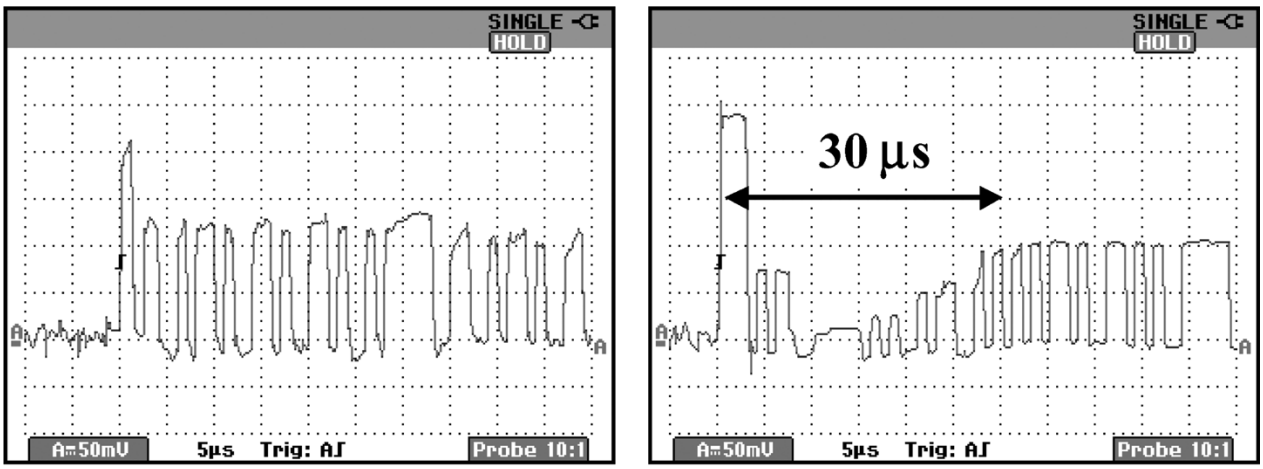

(a) (b)

Fig. 7. (a) Initial AGC recovery of 19-nW average optical power signal with pseudorandom data. (b) Same for 10.1-μW average optical power. In this case, AGC initially takes two coarse steps from 100 to 1 kΩ, followed by several fine steps of increasing PGA gain.

than the LNA, easing the minimum step-size requirement at the charge pump output. This circuit can correct for ambient illumination offsets up to 400 nW with a minimum input-referred step size of 0.8 nW, which is approximately 10 dB below the input referred noise floor. The AOC step size is varied inversely with the PGA gain to maintain constant loop gain, easing the tradeoff between slew rate and loop stability. The worst case optical input-referred offset slew rate is 200 nW/ms.

Manchester coding allows signal peak and minimum levels to be observed in 1.5 symbol periods, after which discrete gain and offset control decisions are made. AOC updates are based solely on the observed "off" level of the OOK signal, which is independent of signal amplitude and is controlled toward zero at the PGA output. AGC updates are based solely on the "on" signal level, which measures the signal amplitude independent of ambient interference once AOC has settled. Both loops operate in a bang/bang manner based on the sign of observed gain and offset errors. The gain control loop automatically switches from fine to coarse control after overflow/underflow of the fine gain control range or when a saturation condition is detected directly at the LNA output. Control updates are applied every 2 bit periods, or every 4 μs at 500 kb/s.

Since the "off" signal level is due to ambient illumination, independent of the received data signal, the AOC loop is able to settle before onset of a new signal. Upon signal onset, only gain control updates are needed to properly condition the signal, so AOC loop speed is not constrained by the 100-μs incident signal recovery requirement. Furthermore, because AGC does not directly affect the measured "off" signal level, the AGC loop can operate much faster than the AOC loop to rapidly acquire new signals, as shown in Fig. 7. In contrast, performing offset control based on the signal mean, directly through ac coupling or though any other linear high-pass filter, requires both gain and offset control adjustments after onset of an OOK signal, resulting in longer acquisition times for new signals.

Worst case recovery over the full dynamic range of the receiver requires 14 gain steps, which are taken in 28 bit periods or 56 μs at 500 kb/s. After recovery of a new signal, AGC logic lengthens the gain control update period to a programmable setting from 128 μs to 4 ms, providing a means to trade off turbulent scintillation tracking rate for resilience to short-link drop outs and residual dither frequency in the AGC loop.

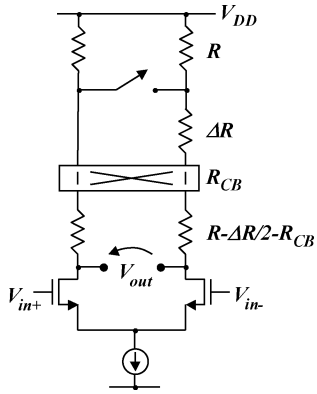

Fig. 8. Single-stage PGA topology with only one of three digital offset trim crossbar switches shown.

Benchtop testing shows that this AGC architecture can accommodate signals up to 21 μW average optical power, limited by the test setup. This is 30.4 dB greater than the minimum detectable optical signal power of 19 nW discussed above and corresponds to a 60.8 dB dynamic range in received photocurrent.

D. Programmable Gain Amplifier

The PGA is implemented as a six-stage amplifier based on the single-stage topology shown in Fig. 8. This is the same topology used for the LNA, with the addition of a single switch used to shunt a portion of the differential load without affecting the common-mode bias condition, providing a single bit of gain control at each PGA stage. As in the LNA, three crossbar switches are used to adjust binary weighted trim resistors for calibration of fixed amplifier offsets, covering an 80 mV output range at each stage. This open-loop topology is acceptable in this application because OOK signals are tolerant of moderate nonlinearity, as long as sufficient bandwidth is maintained to prevent asymmetric rising and falling edge rates. Global bias

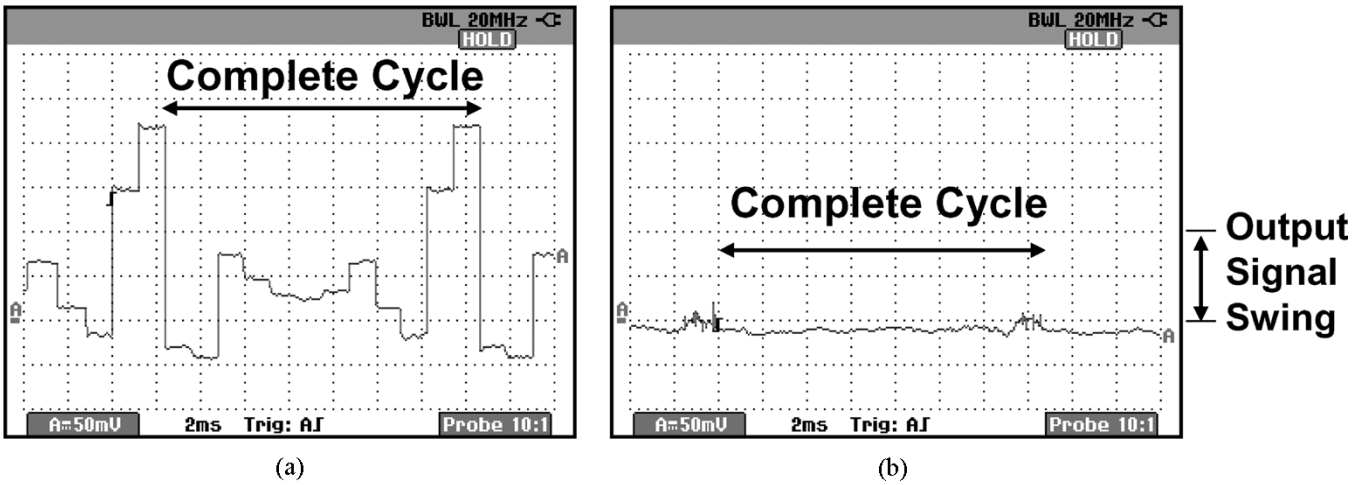

Fig. 9. Measured LNA+PGA output offset voltage as PGA gain is periodically cycled over all 12 gain states from 0 to 33 dB (a) before self-calibration and (b) after self-calibration. The nominal signal range is two vertical divisions at this scale.

adjustment is used to correct for open-loop gain errors due to process variation. In the prototype receiver array, this adjustment is made manually. In principle, this adjustment could be made automatically using a replica bias technique.

The first five PGA stages are programmed to 0 or 6 dB per stage in a thermometer fashion, and the sixth stage is programmed to 0 or 3 dB. Each PGA stage nominally draws 4 μA from a 2.5 V supply. Offset errors at each stage contribute to a total input offset voltage that varies with the gain state. Uncorrected, this error, which is greater than the minimum detectable signal level, leads to undesirable interaction between the AGC and AOC loops. To prevent this, on-chip foreground self-calibration is used to null the overall LNA+PGA input offset voltage across all gain configurations. At each gain setting, the self-calibration routine recursively corrects the offset of each amplifier in the PGA chain, starting at the input stage. Residual uncorrected offset at one amplifier stage is then amplified and included in the offset correction of the succeeding stage. By distributing the offset trim over multiple stages in this manner, residual input referred offsets of less than 100 μV are achieved using coarse trim controls at each stage. This recursive scheme only requires measurement of the sign of each amplifier's output offset voltage, performed by a single comparator from the coarse ADC, as offset correction codes are tested in a cyclic manner. Embedded logic in each pixel uses these comparison results to store the optimal offset correction codes in a 128 bit per-pixel SRAM. These codes are later applied to the PGA at the same time as gain control updates.

Fig. 9 shows the measured output-referred offset voltage of a single LNA+PGA chain over all 12 gain settings before and after self-calibration.

E. Imaging Receiver Array Design

Layout of the single channel optical receiver described above, including necessary circuits for digital clock and data recovery, data framing, and on-chip communication is shown in Fig. 10. This single channel receiver is used as the basis of the 256-element receiver shown in Fig. 11. The total die size is 6.5×6.5 mm^2 with a core imaging area of 28.7 mm^2. Measured

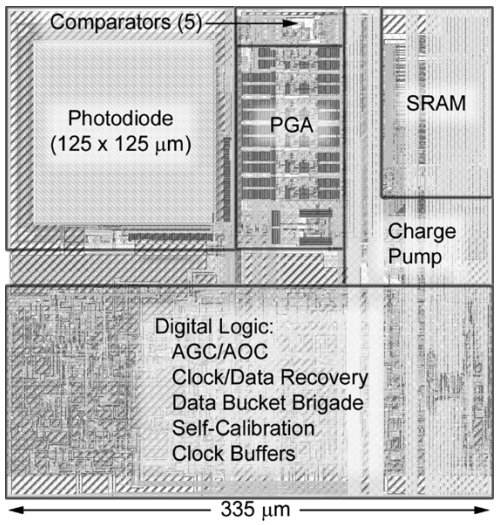

Fig. 10. Layout of single-channel receiver, which is largely covered by a metal light shield.

power consumption for the complete array is 67 mW (56 mW analog, 11 mW digital), excluding I/O.

When a data word up to 16 bit in length is received at any receiver element within the array, it is tagged with a row address and passed from pixel to pixel in a bucket brigade manner toward the bottom edge of the die, where a similar network tags the packet with a column address and passes it toward a shared off-chip bus driver in the lower left corner of the die. In this manner, communication between all receiver channels and the shared data bus is accomplished through nearest neighbor communications, eliminating the need for long on-chip buses and global arbitration networks. Contention for local data passing resources can occur when a receiver element and its upstream neighbor simultaneously try to pass data downstream and at merging data paths at the bottom of each column. Synchronized four-phase handshaking provides arbitration for such cases, passing one word and then the next in a predetermined order without any possibility of metastability. If two merging data streams exceed the maximum data rate of the network, some data words are arbitrarily dropped.

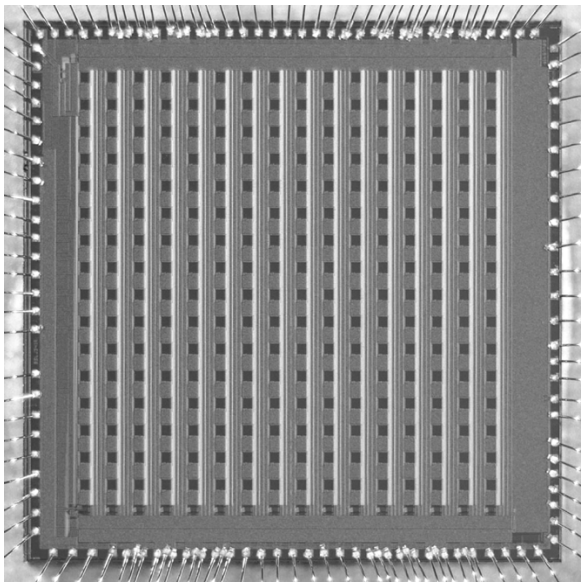

Fig. 11. Microphotograph of the packaged 6.5×6.5 mm^2 256-channel receiver die, based on the single-channel layout shown in Fig. 10. Logic at the bottom of each column receives data words passed down from within the array and passes them in a bucket brigade manner to a shared off-chip bus driver in the lower left corner.

F. Mixed-Signal Integration Challenges

Each element in the imaging receiver incorporates digital logic and SRAM as shown in Fig. 10. The complete receiver array shown in Fig. 11 contains over 1.6 M transistors, most of which operate in the digital domain. In the dual-well process used, the n-well photodiodes connect directly to the bulk substrate, providing a path for digital signals to parasitically couple to the LNA inputs. Furthermore, these sensitive analog nodes are thoroughly integrated throughout the array of digital electronics and are not isolated in a separate portion of the substrate. In addition to the use of an epitaxial wafer with low substrate resistance, circuit techniques are used to minimize substrate noise injection from digital circuits and to provide front-end immunity to any injection that does occur.

First, substrate noise is minimized through design of a custom logic cell library targeted for low energy and area, leading to low slew rates, low supply bounce, and small drain junction areas, all of which reduce substrate noise injection. This is accomplished at the expense of performance by using near minimum size devices, low P/N ratios, and operating core logic circuits from a 1 V supply. The resulting performance penalty is acceptable since the timing requirements for this application are modest compared to intrinsic logic delays in the 0.25 μm CMOS process utilized.

Two circuit techniques shown in Fig. 12 are used to convert any substrate noise at the photodiode into a common-mode signal at the LNA. The first technique is to connect the negative LNA input to the substrate through a replica network that matches the photodiode network, at the expense of doubling the resistor thermal noise power at the LNA input. Capacitor C_D^* is a gate oxide capacitor and not a matched n-well junction diode due to area constraints, leading to a time constant mismatch for the two paths from the substrate to the LNA input.

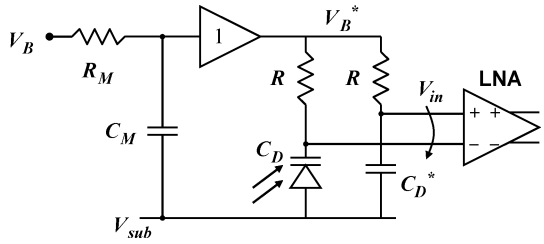

Fig. 12. Substrate noise coupled to the photodiode is converted to a common-mode signal through a replica coupling path and photodiode bias compensation.

This mismatch leads to a bandpass response from the local substrate voltage to the LNA differential input given by

$$\frac{V\text{in}}{V\text{sub}} \approx \frac{s\Delta\tau}{(1+s\tau)^2} \tag{7}$$

where $\Delta\tau$ is the time constant mismatch between the two paths and τ is the nominal time constant of either path. This transfer function has a peak magnitude response of $\Delta\tau/2\tau$ at the signal band edge $\mathrm{s} = 1/\tau$. A manually adjusted 2 bit digital trim of C_D^* is used to correct for gross process variations, matching C_D^* within 10% of C_D and resulting in a maximum substrate noise feed through of -26 dB at the differential LNA input. Because this is a global trim to correct for process variation, local variations in capacitance density degrade this matching.

To further attenuate substrate noise feedthrough to the LNA differential mode input, the substrate voltage is sensed and mirrored to the top of the photodiode bias network as shown in Fig. 12. The unity gain amplifier is based on a single stage plus nMOS source-follower op-amp with an open-loop gain of approximately 40 dB, leading to a closed-loop gain error of approximately 1% while driving the photodiode load. This leads to a band reject response of substrate noise by up to 40 dB across the entire photodiode network, limited by the ac-coupling bandwidth and the op-amp bandwidth of approximately 50 MHz. This response offers complementary rejection to the bandpass response of the replica impedance network discussed above, resulting in broadband attenuation of substrate noise at the differential LNA input.

These substrate noise rejection circuits rely on C_D^* and C_M terminating to the same substrate voltage as C_D within the signal band. This is accomplished by terminating C_D^* and C_M to a substrate tap encircling the photodiode. An epitaxial substrate ensures that the tap ring and the photodiode junction terminate to the same potential through frequencies beyond the signal bandwidth of approximately 2 MHz.

Simulated substrate noise feedthrough of these combined circuits is shown in Fig. 13. In practice, errors such as termination voltage mismatch described above will degrade substrate noise rejection, and this simulation result is optimistic. Although these circuits could not be quantitatively tested because it was not possible to stimulate or probe the substrate independent of the analog ground network in the fabricated prototype, digital crosstalk to the analog signal chain was observed to be below the receiver noise floor measured with all logic disabled.

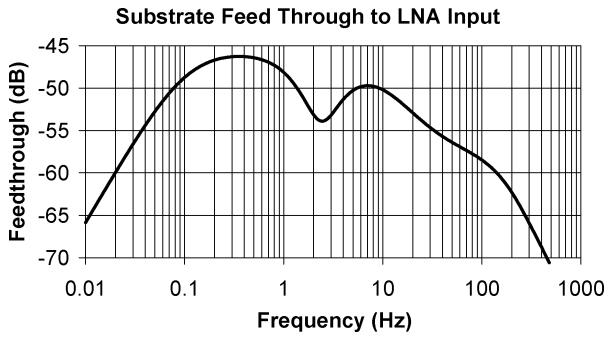

Fig. 13. Simulated differential mode LNA input response to substrate voltage.

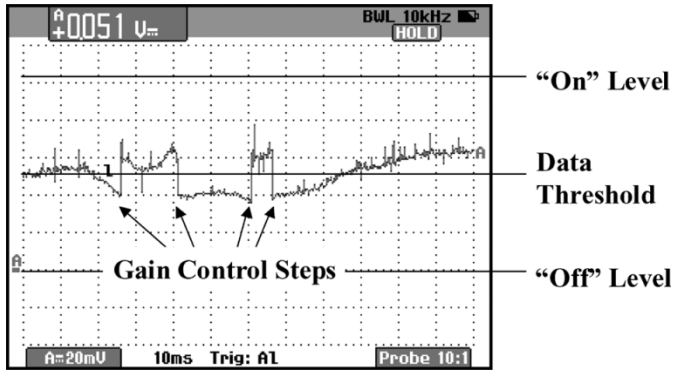

Fig. 14. Mean signal strength observed over 120 ms in a 1.5 km FSO link. AGC makes four discrete gain corrections to keep the signal mean near the threshold used for data slicing.

III. Free Space Link Measurements

Laboratory testing of the fully integrated imaging receiver was performed by directly illuminating individual photodiodes with a manually aligned optical fiber with a known exit beam power. Several dozen receive channels over multiple dies were tested at random without defect, although process variations in C_D and R lead to minor noise performance variations according to (1). Results discussed in the previous section include a minimum detectable signal of 19 nW (-47 dBm), 30.4 dB dynamic range, and recovery of new signals within 56 μs of onset, which is commensurate with the application requirements discussed in the Introduction. To verify the feasibility of this receiver in a real link scenario and to demonstrate the ability of the AGC loop to track and correct for time-varying channel attenuation and ambient illumination, communication over a 1.5 km ground–ground FSO link was tested during both night and daylight hours. Note that stronger turbulence at ground level compared to the target elevation of greater than 100 m reduces the available link range, so this channel can only be qualitatively compared to the target channel in an air–air link and not directly extrapolated.

This 1.5 km FSO link was established in Orinda, CA, after sunset on November 21, 2004, and operated for approximately 36 h. Fixed data words with built-in forward error detection were transmitted at 500 kb/s from a 3 mW average power 660-nm laser with approximately 0.7 mrad beam divergence, similar to that produced by a typical handheld laser pointer. The imaging receiver consisted of the fabricated die positioned in the focal plane of a 25 mm diameter $f/0.72$ aspheric imaging lens, resulting in a receiver field-of-view half angle of 8.4°. A 10-nm bandpass optical filter with 50% in-band transmission was placed before the imaging lens to reduce ambient daylight interference. This entire arrangement was placed on a tripod with coarse angular positioning controls. Fine control of field-of-view orientation was implemented through two-axis micropositioner control of the imaging lens with respect to the imaging array.

While receiving data in this test, the AGC loop was continuously tracking time-varying channel attenuation due to atmospheric turbulence. This is seen in Fig. 14, which shows the mean recovered signal level with AGC operating over a 120 ms interval, measured by low-pass filtering the PGA output at 10 kHz. During this interval, four discrete gain control steps are taken at the indicated times to maintain a mean signal level near the data threshold.

Over the course of the 36 h test, communication was only intermittently available. During both nighttime and daylight hours, throughout the test there were intervals greater than 30 min with continuous communication during which more than 99% of data words were received without bit errors over any 10 s subinterval. This error rate should be acceptable in the intended visual surveillance application, where minor data degradation is tolerable. Tolerance to single bit and burst errors in an application of course depends on data coding, potential use of error detecting/correcting codes, and the transmission control protocol, which are not discussed here. Link dropouts over the course of the test are presumed to be caused by fog, haze, and beam wander due to slowly varying large-scale temperature gradients in the atmosphere, all of which were directly observed at some points during the free-space link test. In the target application, two transceivers with bidirectional communication would use active beam steering to correct for beam wander, as discussed in [3].

IV. Concluding Remarks

A 256-element imaging receiver for FSO communication has been integrated in a 0.25 μm CMOS die measuring 6.5×6.5 mm^2. The 1.6 M transistor mixed-signal circuit has a measured optical sensitivity of -47 dBm and a 30.4 dB dynamic range at 500 kb/s with Manchester data coding. Total power consumption for the die is 67 mW, excluding I/O. Communication over a 1.5 km horizontal path near ground level was demonstrated with a 3 mW, 0.7 mrad transmission beam at 660 nm and a 25 mm receiver aperture. This demonstrates the feasibility of long-range FSO communication with small low-power transceivers, despite the presence of time-varying channel attenuation due to atmospheric turbulence. Noise analysis shows that, despite arbitrarily low resistor noise current, the commonly used transimpedance amplifier topology can only achieve significantly better noise performance than a series sense impedance if the core MOS voltage amplifier has $g_m \gg 2\pi BC$. This admittance is dominated by the application-dependent signal bandwidth B and photodiode capacitance C_D. Circuit techniques have been presented to cope with the mixed-signal challenge of sensing single-ended, bulk substrate coupled signals throughout a large digital circuit.

REFERENCES

[1] P. Djahani and J. M. Kahn, "Analysis of infrared wireless links employing multi-beam transmitters and imaging diversity receivers," *IEEE Trans. Commun.*, vol. 48, no. 12, pp. 2077–2088, Dec. 2000.

[2] M. Last, B. Fisher, C. Ezekwe, S. Hubert, S. Patel, S. Hollar, B. S. Leibowitz, and K. S. J. Pister, "Video semaphore decoding for free-space optical communication," *Proc. SPIE*, vol. 4303, pp. 148–158, 2001.

[3] M. Last, B. S. Leibowitz, B. Cagdaser, A. Jog, L. Zhou, B. E. Boser, and K. S. J. Pister, "Toward a wireless optical communication link between two small unmanned aerial vehicles," in *Proc. IEEE Int. Symp. Circuits and Systems*, vol. 3, May 2003, pp. 930–933.

[4] B. P. Venet, "Optical scintillometry over a long, elevated horizontal path," *Proc. SPIE*, vol. 3381, pp. 212–219, 1998.

[5] L. C. Andrews and R. L. Phillips, *Laser Beam Propagation Through Random Media*. Bellingham, WA: SPIE Opt. Eng. Press, 1998.

[6] P. Popescu, A. Solheim, and M. Wight, "Experimental monolithic high speed transceiver for Manchester encoded data," in *Proc. Bipolar/BiCMOS Circuits and Tech. Meeting*, Minneapolis, MN, Oct. 1995, pp. 110–113.

[7] D. J. Heatley, "Optical receivers," in *Design of Analog VLSI Circuits for Telecommunications and Signal Processing*, Y. E. Franca and Y. Tsividis, Eds. Englewood Cliffs, NJ: Prentice-Hall, 1993.

[8] D. B. Schwartz, R. E. Howard, and W. E. Hubbard, "A programmable analog neural network chip," *IEEE J. Solid-State Circuits*, vol. 24, no. 2, pp. 313–319, Apr. 1989.

Brian S. Leibowitz (S'96–M'98) received the B.S. degree in electrical engineering from Columbia University, New York, NY, in 1998 and the Ph.D. degree from the University of California, Berkeley, in 2004. His graduate research on CMOS imaging receivers was supported by a generous fellowship from the Fannie and John Hertz Foundation.

Since 2004, he has been with Rambus, Inc., Los Altos, CA, where he has been involved with the development of high-speed serial transciever circuits. His primary research interests are in analog and mixed-signal integrated circuits.

Bernhard E. Boser (S'79–M'83–SM'01–F'03) received the Diploma in electrical engineering from the Swiss Federal Institute of Technology, Zurich, in 1984 and the M.S. and Ph.D. degrees from Stanford University, Stanford, CA, in 1985 and 1988, respectively.

From 1988 to 1992, he was a Member of Technical Staff with the Adaptive Systems Department, AT&T Bell Laboratories. In 1992, he joined the faculty of the Department of Electrical Engineering and Computer Science, University of California, Berkeley, where he also serves as a Director of the Berkeley Sensor and Actuator Center. His research is in the area of analog and mixed-signal circuits, with special emphasis on analog–digital interface circuits and micromechanical sensors and actuators.

Dr. Boser has served on the program committees of the IEEE International Solid-State Circuits Conference, the Transducers Conference, the VLSI Symposium, and the Sensor and Actuator Workshop. He has served as the Editor-in-Chief of the IEEE JOURNAL OF SOLID-STATE CIRCUITS.

Kristofer S. J. Pister received the B.A. degree in applied physics from the University of California at San Diego, La Jolla, in 1986, and the M.S. and Ph.D. degrees in electrical engineering from the University of California, Berkeley (UC Berkeley), in 1989 and 1992. respectively.

From 1992 to 1997, he was an Assistant Professor of electrical engineering with the University of California, Los Angeles, where he helped developed the graduate MEMS curriculum and coined the phrase Smart Dust. Since 1996, he has been a Professor of electrical engineering and computer sciences with UC Berkeley. In 2003 and 2004, he was on leave from UC Berkeley as CEO and then CTO of Dust Networks, Berkeley, CA, a company he founded to commerciallize wireless sensor networks. He has participated in many government science and technology programs, including the DARPA ISAT and Defense Science Study Groups, and he is currently a member of the Jasons. His research interests include MEMS, microrobotics, and low-power circuits.