Integrated Sensor and Range-Finding Analog Signal Processor

Andrew Gruss, L. Richard Carley, *Member, IEEE*, and Takeo Kanade, *Senior Member, IEEE*

Abstract —In this paper we present an IC array of photosensor and analog signal processor cells that acquires 1000 frames of light-stripe range data per second—two orders of magnitude faster than conventional light-stripe range-finding methods. The highly parallel range-finding algorithm employed requires that the output of each photosensor site be continuously monitored. Integration of processing at the point of sensing makes implementation of this algorithm practical. Prototype high-speed range-finding systems have been built using a 5×5 array and a 28×32 array of these sensing elements.

I. INTRODUCTION

RANGE FINDING, or measurement of the three-dimensional profile of an object or scene, is a critical component for many robotic applications. Numerous range-finding techniques have been developed [1], and of these light-stripe range finding is one of the most common and reliable methods. A conventional light-stripe range finder operates in a step-and-repeat manner: a light stripe is projected onto a scene, a video image is acquired, the position of the projected stripe in the image is extracted, the stripe position is stepped, and the process repeats until the entire scene has been scanned. The rate at which frames of range data can be acquired using this method is limited by the time needed to acquire and process each video image. Note that the number of video images required to build a complete range map increases linearly with the desired spatial range resolution. Typically, on the order of one second is required to acquire a complete range image in this manner.

The integrated range finder is based on a modification of the basic light-stripe range-finding technique described by Sato *et al.* [2] and Kida *et al.* [3]. In the modified method, parallel light-stripe range finding, the range image is constructed in a parallel manner rather than a serial one. The light stripe is swept continuously across the scene and the time at which the image of the stripe crosses each cell of the range sensor indicates the range along that photodiode's line of sight. In this manner, an entire range map is acquired in parallel during a single continuous sweep of the stripe across the scene and the total time of acquisition is independent of the range map resolution. If the geometry of the optics and the speed of the light stripe remain fixed, increasing the size of the sensor array just increases the visual angle that is imaged. Thus, the two orders of magnitude speed increase achieved by the small prototype array presented in this paper will not be decreased as the number of cells in the sensor array increases.

The parallel range-finding algorithm requires the detection of the light stripe passing across the photosensor, which generates a narrow pulse of current at the sensor's output. Accurate detection of this narrow pulse requires a high-bandwidth connection between the photosensor and the signal-processor that detects the time of the peak. Sato *et al.* [2] constructed a prototype system to demonstrate the parallel range-finding method using a small array of discrete photosensors connected to separate peak detection hardware. For large arrays, the only practical way to provide a high-bandwidth link between each photosensor and the signal-processing electronics is to integrate both photosensors and signal-processing circuitry into a dense array of range sensing cells. In order to keep the cell area small, the necessary signal-processing functions have been implemented using primarily analog circuitry.

In this paper we present experimental results from an array of cells, each of which contains a photodiode and the analog signal-processing circuitry needed for light-stripe range finding. Prototype circuits were fabricated through MOSIS [4] in a 2-μm CMOS p-well double-metal, double-poly process. This design builds on some of the ideas that have been developed for IC's that integrate signal-processing circuitry with photosensors, e.g., an optical mouse [5], an artificial retina [6], and an optical position encoder [7]. In the case of light-stripe range finding, the increase in cell complexity from sensing only to sensing and processing makes the modification of the operational principle of range finding practical, which in turn results in a dramatic improvement in performance.

Section II provides a more complete description of the conventional and parallel light-stripe range-finding methods. Section III describes the design of a single cell of the prototype range sensor IC, which contains a photosensor and an analog signal processor. The architecture of the prototype range sensor IC and the light-stripe range-finding system is also discussed. Section IV presents experimental measurements of the prototype range-finding system and Section V concludes the paper.

A. Gruss and T. Kanade are with the School of Computer Science, Carnegie Mellon University, Pittsburgh, PA 15213-3890.
L. R. Carley is with the Department of Electrical and Computer Engineering, Carnegie Mellon University, Pittsburgh, PA 15213-3890.

Fig. 1. Light-stripe range finding.


Fig. 2. 2D array of smart photosensors.

Fig. 3. Parallel light-stripe technique.

Fig. 4. Measuring range information.

II. Light-Stripe Range Finding

The conventional step-and-repeat light-stripe range-finding procedure is presented first in this section for comparison with the parallel method on which the IC sensor is based. Fig. 1 shows the geometric principle on which a light-stripe range finder operates. The scene is illuminated with a vertical plane of light. The light is intercepted by an object surface in the path of the beam and, when seen by a video camera placed to the left of the light source, the plane of light appears in the image as a stripe which follows the surface contour of objects in the scene.

Range data along the contour can be calculated using the principle of triangulation. In Fig. 1, the equation for the position of the plane of light is known because the projection angle is controlled by the mirror positioning assembly. The line of sight for each point p on the image of the stripe can also be determined by tracing a line from the camera focal point through p. The intersection of the line-of-sight ray with the projected plane of light uniquely determines the three-dimensional position of that point on the object's surface.

A. Conventional Light-Stripe Range Finding

A conventional light-stripe range finder collects range data for an entire scene sequentially—iterating the process of fixing the stripe on the scene, taking a picture, and signal processing each picture to detect the position of the light stripe until the entire scene has been scanned. Though practical, the maximum rate at which frames of range data can be acquired by this technique is severely limited because the number of pictures which must be taken and processed increases linearly with the desired range resolution. Assume that a video camera image has N rows. Typically, N ranges between 256 and 512 samples and the time required to acquire and process an image frame ranges between 1/30 and 1/10 s. Thus, camera-based range finders typically require between 8.5 and 51 s. Through the use of high-speed cameras and sophisticated optics and electronics, the maximum speed of conventional light-stripe range finders has been increased to ten frames of range data per second.

B. Parallel Light-Stripe Range Finding

In the parallel range-finding technique, the video camera is replaced by a two-dimensional array of smart photosensitive cells as illustrated in Fig. 2. Range data are not acquired in a step-and-repeat manner. Instead, the plane of light is swept once across the scene from left to right at a constant angular velocity. This slight modification of the conventional light-stripe range-finding algorithm transforms the range-

finding process from a sequential one into a parallel one, and results in a qualitative change in the way the range measurement is determined.

Referring to Fig. 3, each cell S_m detects and remembers the time t_m at which it observed the peak incident light intensity $I_m(t)$ during each sweep of the stripe. Observe that each cell predefines a unique line of sight and that the time information t_m recorded by the cell determines a particular orientation of the stripe $\theta(t_m)$. Recalling the geometry in Fig. 1, one sees that this information is sufficient to calculate the three-dimensional position of the imaged object point, again using triangulation. The data gathered during one pass of the stripe in each cell of a $M \times M$ array of these smart sensing elements are sufficient to calculate the $M \times M$ range map of an imaged scene. For the new method, the sweep time T_s is only limited by the photoreceptor sensitivity at the chosen bandwidth. The resolution (the number of cells) is limited by integration technology, that is, the size of the sensor and the necessary signal-processing circuitry and the yield characteristics of the available technology.

C. Accuracy of the Range Measurement

The accuracy of the range measurement is actually improved when using the parallel light-stripe range-finding method. As illustrated in Fig. 4, the accuracy of the conventional camera-based method is limited by the number of pixels in the horizontal scan line (typically 256 to 512 for a camera) because the range information is derived from the position of the peak of the light intensity profile on the scan line. Subpixel peak localization techniques, which can at best provide one tenth of a pixel spacing accuracy, are often employed to work around this problem. It is interesting to note that since the accuracy relies on the interpolation of a spatially sampled profile, an extremely fine light stripe which

Fig. 5.   Sensing element circuitry.

Fig. 6.   Photodiode area and pulse shape.

provides good $x-y$ resolution is not necessarily the best for obtaining z accuracy. However, in the parallel range-finding method the peak of the continuous uninterpolated time profile of intensity from the same cell is detected in the time domain with much greater accuracy. The spatial position of the light stripe is accurately determined by a shaft encoder on the light-stripe projector; hence, a fine light stripe directly contributes to an increase of both $x-y$ resolution and z accuracy.

In the standard technique, range data are extracted by measuring intensity changes as a function of position in the sensor data (video images). The parallel technique monitors intensity changes as a function of time on line-of-sight rays fixed in space. Thus, the parallel technique is better suited to extract range data from an actively illuminated scene.

III. IC RANGE SENSOR

Functionally, each element of the smart photosensitive array converts light energy into an analog voltage, determines the time at which the voltage peaks, and remembers the time at which the peak occurred. The implementation of these functions requires that the photoreceptor and signal-processing circuitry be integrated into each cell. The photosensor design must take into consideration trade-offs in cell size, power dissipation, bandwidth, sensitivity, and accuracy. In order to minimize cell die area, the current cell design implements each of the necessary functions in the simplest possible manner (see Fig. 5), while still providing the necessary functionality. A photodiode is used rather than a phototransistor in order to maximize the bandwidth. To detect the time at which the light stripe crosses the photodiode, a threshold detector follows the preamp. Although some form of filtering between the preamp and the comparator would be desirable, its implementation in a very small die area is problematic. Storage of the time at which the light stripe is detected is implemented in a straightforward manner. An analog voltage ramp, phase locked to the motion of the stripe (the voltage is proportional to the angular stripe position), is distributed to every cell. A track-and-hold circuit (one switch and one capacitor) within each cell tracks the analog ramp until light falling on the cell causes the photodiode signal to exceed a global threshold voltage. The value of the ramp voltage held on the track-and hold capacitor when the threshold circuit trips is thus proportional to the time at which the cell observes a change in intensity caused by the light stripe. As described above, these held time stamps are the raw data from which range information is derived. The light-stripe range-finding operation is divided into two phases, data collection and data readout. Phase one begins by resetting all of the cells, and then each cell detects the passing of the image of the light stripe as it sweeps across the sensor IC. During phase two, the cell data are read out from the

chip in a raster scan fashion. The chip is then reset and another scan begins.

A. Integrated Photodiodes and Preamplifier

Photodiodes are critical to the sensitivity and bandwidth of the sensor cells. The current output from the photodiode, at a given incident light intensity, and the junction capacitance are directly proportional to the photodiode area. The more area devoted to photodiode structures the better the optical sensitivity. In the prototype array, the photodiodes are approximately 8000 μm^2 in area, one fourth of the total area budgeted for a cell. In a CMOS process, one of the highest sensitivity photodiode structures is the well–substrate junction [7]. In a p-well CMOS process, this vertical photodiode structure is constructed using the n-type substrate as the cathode and the p-type well as the anode. An additional p^+ implant is driven into the well to reduce the surface resistivity of the anode to which contact is made. Photodiodes with and without a passivation layer were fabricated; however, the passivation layer caused only a slight degradation of the sensitivity at the stripe wavelength.

The shape of the photodiode area is important. Consider the image of a stripe moving across the photodiode as sketched in Fig. 6. The amount of light collected varies because the length of the intersection between the stripe and the photodiode area varies. In the case of a square photocell, the output pulse shape depends on stripe orientation; for a vertical stripe, the output pulse observed will be square, but a 45° stripe will produce a triangle wave. This difference in pulse shapes translates to an error in estimating the time at which the peak of the stripe occurs. Therefore, the photodiodes are circular so that the pulse shape will be independent of stripe orientation.

Finally, the photodiode structures are surrounded with guard rings to help prevent noise currents introduced into the substrate at other points from contaminating the photocurrent, and to prevent photocurrent-induced latch-up at high light levels. Photodiode anodes are high-impedance nodes and will be susceptible to noise coupling from other chip circuitry. In particular, digital clock transitions occurring while stripe detection is taking place could inject noise and corrupt the delicate photocurrent measurements. We considered using a switched-capacitor filter [8] in the prototype cell design between the preamp and the threshold detector, because switched-capacitor filters can be constructed in a small die area. However, even though techniques exist which can help alleviate problems [9] associated with mixed analog–digital IC design, uncertainty about the amount of substrate noise injected by high-frequency clocks led us to prefer continuous-time signal-processing circuits. While photodiode outputs are being monitored, digital transitions occur only when a given cell detects the light stripe. However, since this firing of the latches within each cell occurs asynchronously throughout the array, these transitions

Fig. 7. Transimpedance preamplifier.

Left-to-right / Stripe Scan   Right-to-left / Stripe Scan

Fig. 8. Amplified photocurrent oscilloscope waveforms.

cause only a small amount of local disturbance. In the prototype, digital clocks run only during the readout phase, when photodiode outputs are not being monitored.

In addition to the light-stripe range-finding IC arrays, a test IC that included just photodiodes and the transimpedance preamp circuit (see Fig. 7) was also fabricated using the same IC process. In this circuit the preamp output was brought to pads so that the performance of the photodiode and preamp could be characterized. The circuit of Fig. 7 operates as a logarithmic transimpedance amplifier in the following manner. Transistor Q_1's drain current is equal to that of the current (photocurrent plus leakage) produced by the reverse-biased photodiode. Subthreshold operation of transistor Q_1 is thus insured by this small (typically less than a nanoampere) drain current. The second PMOS transistor, Q_2, acts as a common-source amplifier stage. Feedback tends to keep the voltage at transistor Q_1's source terminal equal to Q_2's V_{GS}—independent of light level because the drain current of transistor Q_2 is fixed by the load device Q_3. The bandwidth of the transimpedance amplifier is also enhanced because the voltage swing seen across the photodiode's junction capacitance is attenuated by the feedback action across transistor Q_2. Changes in the drain current of transistor Q_1, operating in subthreshold, produce logarithmic changes in its gate-to-source voltage [10]. Thus changes in the output voltage of the transimpedance amplifier, the gate of transistor Q_1, will vary as the log of the current produced by the photodiode.

The logarithmic characteristic causes the effective gain of the preamplifier to be inversely proportional to the sum of the photocurrent and photodiode leakage current. Similarly, the bandwidth of this circuit is directly proportional to this same current. In practice, a uniform background illumination of the sensor IC can be used to make a trade-off between sensitivity and bandwidth. Conversely, illumination from the scene can generate a nonuniform gain across the photodiode/preamp array. In order to remove nonuniform background illumination, a wavelength selective optical filter, with a center wavelength at that of the laser generating the light stripe, is placed in the optical path of the sensor.

The preamp test chip was mounted in the film plane of a standard 35-mm camera and a moving stripe was used to generate light stimuli. Fig. 8 shows the experimental results of applying a swept infrared laser stripe across two adjacent cells. It demonstrates that the signal has a clear peak whose location moves as the distance to the surface changes. The bandwidth of these test sensors is compromised by the connection to a bonding pad. However, the ability of this photodiode/preamplifier combination to successfully detect stripes swept at rates corresponding to 1000 range frames/second has been experimentally observed for relatively reflective surfaces (higher light input signal level). The maximum rate for dark surfaces drops to 100 frames/second.

B. Stripe Detection and Time Storage

In the sensor array prototype, light-stripe peak detection is done by thresholding. The amplified photodiode signal is compared to a reference voltage, ranging between 2.5 and 3.0 V, set to produce an output when the stripe image passes across the photodiode. The reference input is directly connected to an off-chip voltage source through one of the bonding pads. Detection of the stripe trips the RS flip-flop in the cell, indicating that the light-stripe has been detected and that the time-stamp value should be held (see Fig. 5).

The decision to avoid having any digital signals switching during the sweep directly affected the comparator design. The comparator in each sensing element is nonregenerative —essentially an uncompensated two-stage op-amp circuit [11] with no applied external feedback and no compensation. Although a clocked cell design would allow the use of regenerative comparator topologies, which are generally faster for a given amount of die area, the potential danger of injecting clock noise into the photocurrent measurement outweighed the potential benefits.

A global time signal indicating the position of the light stripe is broadcast to each cell from an off-chip reference as a 0–5-V analog ramp synchronized with the stripe rotation. Sweeps (ramp duration) of down to 1 ms were employed. The time-stamp voltage is held on the capacitor of a track-and-hold (T/H) circuit when the latch in the cell is tripped by the detection of the stripe. Representing time in an analog form has area advantages over the digital equivalent of latching the value of a continuously running counter and, in addition, avoids the problems associated with mixing sensitive analog circuitry with digital logic as previously described. Unfortunately, information stored in analog form is subject to corruption from sources such as switch charge injection and voltage droop. However, with careful circuit design, errors due to charge injection, for example, can be kept to about 0.5%, yielding acceptable range accuracy. In addition, any systematic components of these errors can be calibrated and subtracted after the range data have been transmitted to a digital computer. In that case, the primary sources of error are random noise sources and nonlinearities.

The primary reason for choosing analog storage over digital was die area. In a digital implementation, the multibit (at least 8 b) digital time-stamp signal bused over the entire chip and the multibit digital latch within each cell require a large amount of the die area. An analog time-stamp voltage can be

Fig. 9.   Layout of the light-stripe cell.

Fig. 10.   Off-chip readout circuitry.

Fig. 11.   6×6 sensor array chip.

circuit. Because the voltages held on the T/H capacitors will tend to droop due to substrate leakage currents, it is important to off-load range samples immediately after the data acquisition phase of operation has been completed. An external 12-b 1-Msample/s analog-to-digital (A/D) converter is used to convert the multiplexed analog time values into digital form. Simple processing to convert this raw time-stamp data into 3D (x, y, z) range map values and correction for individual cell offset and scale factor errors will be performed by a host computer system.

D. Sensor IC Layout

A die photo of one prototype sensor chip, including an array of 6×6 sensing elements and readout circuitry, is shown in Fig. 11. Each cell (sensor and signal processor) is roughly 250 μm×250 μm. As mentioned previously, the photodiode areas are circular to provide consistent output pulse shapes. The T/H capacitor has been placed in the corners around photodiodes to minimize wasted die area. Sensing element processing and support circuitry is sandwiched between the photodiodes. Column-wise readout circuitry, the output op amp, and integrating capacitor are seen in a row just above the top of the array. As is usually the case with photosensitive IC sensors (e.g., CCD's), second-level metal is used as a light shield over every active device (except the photodiode) to prevent the spurious generation of photocurrents, or photoinduced latch-up.

broadcast over the entire chip on a single wire and the circuitry to record an analog time value consists only of a holding capacitor and a switch.

In the prototype, a double-poly process was used to provide high-quality linear capacitors that exhibit good matching across the die. The matching of these capacitors across the sensor chip will in large part determine the variance in voltage reported by individual sensors for given time values. However, this is not a fundamental limitation as variations between capacitors would also be taken out during the calibration procedure and corrected digitally off chip. In addition, the area requirements of the capacitor are less than one might expect. Although the 1-pF capacitor used would need to be a rectangle 45 μm on a side, the actual implementation uses the space around the circular photodiode, which would be hard to make use of otherwise (see Fig. 9).

C. Stored Time Readout

When a scan has completed, the hold capacitor in each cell of the array is at the voltage corresponding to the time at which the stripe was detected along its line of sight. Readout of these stored range values proceeds in a raster scan fashion, like a standard CCD camera except that it is range, not intensity, which is being read out. A shift register, wound through the array, gates the charge held on each cell's T/H capacitor sequentially onto a single analog readout bus. The accumulated charge is integrated by an on-chip op-amp integrator and presented to an output pin as a buffered output voltage using the circuit shown in Fig. 10. The output voltage presented is halved to avoid headroom problems in the output op amp and inverted in polarity by the integration

E. Test System

A prototype range-finding system has been built that uses the integrated sensor/signal processor IC described above. Essential components in this system include stripe-generation hardware, range sensor IC and support electronics, range sensor optics, and a host interface. The range sensor IC was mounted in the film plane of a 35-mm camera body, which provided a convenient mechanism for incorporating focusing optics and for sighting the range finder. The stripe was generated using a 10-mW infrared (780 nm) laser-diode-based projector. A mirror mounted on a galvonometer provided the means of sweeping the stripe. Two scans (one from left to right, the other from right to left) were generated during each period of the 500-Hz triangle wave used to drive the galvonometer, providing the system with a maximum of

Fig. 12. Range-finding experiment geometry.

Fig. 13. 5×5 range-finder operation—inclined plane at 50 cm.

1000 sweeps per second. The experimental system has a 50-cm baseline between the scan generator and the sensor.

IV. Experimental Results

Figs. 13–15 show waveforms observed during the operation of a 5×5 range sensor prototype. In each of the figures, the top trace shows the output of the op-amp integrator that reflects the charge stored on each of the cell hold capacitors. The five groups of five pulses that appear are the output from the cells of a 5×5 sensor array being read out in raster-scan fashion as previously explained. In order to better understand the waveforms shown, refer to Fig. 12. In the figure, a single sensor's view of two object positions, O_{near} and O_{far}, is outlined. For the stripe scan direction shown, an object point at O_{near} will generate a range angle, θ_{near}, greater than the range angle measured, θ_{far}, for an object point O_{far}. The time-stamp voltage value of an element in the array, as measured at the output of the on-chip op-amp integrator, is proportional to the angles shown in the figure. Remember that the readout integrator inverts the voltage corresponding to the time the stripe was detected. Thus, low pulses (below the 2.5-V idle value) represent stripe events recorded late in the sweep of the stripe (and close in range). High-valued pulses represent early range events, from objects positioned further from the sensor. The baseline distance (stripe axis of rotation to sensor focal point) for the three experiments is 50 cm.

The second trace shows the analog time-stamp reference ramp being broadcast to all elements during the range data acquisition phase, at a rate corresponding to approximately 1000 frames/second. The time-stamp input idles at a constant 0-V value while range data readout is taking place. Data readout occurs after completion of a scan. The bottom trace is an index pulse generated by the stripe generation assembly that indicates the stripe is positioned at the beginning ($t = 0$) position. Fig. 13 shows a posterboard (plane) held approximately 50 cm from the sensor. Moving the posterboard to a distance of 90 cm changes the output to that shown in Fig. 14. With the poster board still at 90 cm, a pencil held vertically in the field of view of the sensor at a distance of 55 cm produces the output shown in Fig. 15.

Fig. 17 shows recent results obtained from a 28×32 element range sensor array (Fig. 16). As in the previous oscilloscope photos, the output waveform shown is that of the range data raster produced during the readout phase. The acquisition phase for this data occurred just before the output shown at a sweep rate corresponding to 100 frames/second. As before, low voltage represents a "range

Fig. 14. 5×5 range-finder operation—inclined plane at 90 cm.

Fig. 15. 5×5 range-finder operation—vertical pencil at 55 cm.

event" occurring late in time (e.g., values close to the top of the photograph are early range events). The nearly vertical lines are range samples from the 32 chip columns, too compressed in time to be individually visible as was the case for the 5×5 data. The fact that a 2D image can be seen in Fig. 17 is an artifact of the way the pieces of the range-finding system have been arranged. The sensor's columns are oriented to be parallel to the left-to-right sweep of the stripe. Thus, the sensor systems x axis (left to right) and z axis (distance) both map to the y axis of the oscilloscope. The scope's x axis (time) is where the y axis (up and down) of the sensor system is mapped.

V. Conclusion

Advances in VLSI technology have made it possible to create smart sensors, ones that integrate sensing and signal processing. Analog signal processing is well suited for performing the kinds of computations necessary to achieve

Fig. 16.   28×32 sensor array chip.

Fig. 17.   28×32 range-finder operation—coffee cup at 24 cm.

intelligent sensors. The precision required of the signal-processing operations is relatively low. The integrated circuitry which performs these computations on analog signals can be quite small in area, allowing a higher number of processors to be integrated on a single sensor IC. Analog processing circuits need not generate the switching noise associated with digital implementations. Such noise is a factor when working with the low-level outputs from an on chip sensor. Finally, sensed data are by their nature analog. Analog processing avoids the need to convert sensed data before computations can take place. This is particularly important when sampling causes information loss due to aliasing in the frequency domain.

We have demonstrated a VLSI smart sensor which performs light-stripe range finding and is capable of acquiring 100 to 1000 frames of range data per second. One of the most distinguishing features of this approach is that it is not just a parallel implementation of known algorithms using

VLSI technology to achieve increased speed. Rather, it demonstrates that integration of sensing and signal processing can make it possible to modify the operating principles of the information acquisition method (in our case, light-stripe range finding) which can result in a qualitative improvement in performance.

REFERENCES

[1] P. J. Besl, "Range imaging sensors," General Motors Research Labs., Warren, MI, Res. Publ. GMR-6090, Mar. 1988.
[2] Y. Sato and K. Araki, and S. Parthasarathy, "High speed rangefinder," in Optics, Ilumination, and Image Sensing for Machine Vision II, SPIE, vol. 850, pp. 184–188, 1987.
[3] T. Kida and K. Sato, and S. Inokuchi, "Realtime range imaging sensor," in Proc. 5th Sensing Forum, Apr. 1988, pp. 91–95 (in Japanese).
[4] G. Lewicki et al., USC MOSIS User's Manual, Inform. Sci. Inst., Univ. Southern Calif., Marina Del Rey, CA, 1988.
[5] R. F. Lyon, "The optical mouse, and an architectural methodology for smart digital sensors," Xerox Palo Alto Res. Center, Palo Alto, CA, Tech. Rep. VLSI-81-1, Aug. 1981.
[6] M. A. Sivilotti, M. A. Mahowald, and C. A. Mead, "Real-time visual computations using analog CMOS processing arrays," in Advanced Res. VLSI — Proc. 1987 Stanford Conf., P. Losleben, Ed. Cambridge, MA: MIT Press, pp. 295–312.
[7] P. Aubert and H. Oguey, "An application specific integrated circuit (ASIC) with CMOS-compatible light sensors for an optical position encoder," in Proc. IEEE Custom Integrated Circuits Conf., May 1987, pp. 712–716.
[8] P. E. Allen and E. Sanchez-Sinencio, Switched Capacitor Circuits. New York: Van Nostrand Reinhold, 1984.
[9] J. A. Olmstead and S. Vulih, "Noise problems in mixed analog-digital integrated circuits," in Proc. IEEE Custom Integrated Circuits Conf., May 1987, pp. 659–662.
[10] C. Mead, Analog VLSI and Neural Systems. Reading, MA: Addison-Wesley, 1989.
[11] R. Gregorian and G. C. Temes, Analog MOS Integrated Circuits for Signal Processing. New York: Wiley, 1986.

Andrew Gruss was born in Latrobe, PA, in 1957. He received the B.S. and M.S. degrees in electrical engineering from Carnegie Mellon University, Pittsburgh, PA, in 1979 and 1981, respectively. He is currently working towards the Ph.D. degree at Carnegie Mellon, exploring the use of integrated sensing and processing techniques in robotic sensing applications.

In 1979 he joined the research staff in the School of Computer Science at the university, working in the areas of digital systems design, computer networking, and data acquisition system design.

L. Richard Carley (S'77-M'84) received the S.B. degree from the Massachusetts Institute of Technology (MIT), Cambridge, in 1976 and was awarded the Guillemin Prize for the best EE undergraduate thesis. He remained at MIT where he received the M.S. degree in 1978 and the Ph.D. degree in 1984.

He has worked for MIT's Lincoln Laboratories and has acted as a consultant in the area of analog circuit design and design automation for Analog Devices and Hughes Aircraft among others. In 1984 he joined Carnegie Mellon University, Pittsburgh PA, where he is currently an Associate Professor of Electrical and Computer Engineering. His research is in the area of analysis, design, automatic synthesis, and simulation of mixed analog/digital systems.

Dr. Carley received a National Science Foundation Presidential Young Investigator Award in 1985, and a Best Paper Award at the 1987 Design Automation Conference.

Takeo Kanade (M'80–SM'88) received the Ph.D. degree in information science from Kyoto University, Japan, in 1974.

He is a Professor of Computer Science and the Co-Director of the Robotics Institute at Carnegie Mellon University (CMU), Pittsburgh, PA. He has worked on various problems in vision, sensors, manipulators, and mobile robots. Currently he is engaged in three robotics research programs at CMU as the Principal Investigator: image understanding, the autonomous land vehicle (NAVLAB) vision system, and the NASA Mars rover. He is also the Chairman of the newly established robotics Ph.D. program at CMU.