

Smart Dust: Communicating with a Cubic- Millimeter Computer



The Smart Dust project is probing microfabrication technology's limitations to determine whether an autonomous sensing, computing, and communication system can be packed into a cubic-millimeter mote to form the basis of integrated, massively distributed sensor networks.

*Brett
Warneke
Matt Last
Brian
Liebowitz
Kristofer S.J.
Pister*
University of
California,
Berkeley

Decreasing computing device size, increased connectivity, and enhanced interaction with the physical world have characterized computing's history. Recently, the popularity of small computing devices, such as handheld computers and cell phones, burgeoning Internet growth, and the diminishing size and cost of sensors—and especially transistors—have accelerated these trends. The emergence of small computing elements, with sporadic connectivity and increased interaction with the environment, provides enriched opportunities to reshape interactions between people and computers and spur ubiquitous computing research.¹

The Smart Dust project² is exploring whether an autonomous sensing, computing, and communication system can be packed into a cubic-millimeter mote (a small particle or speck) to form the basis of integrated, massively distributed sensor networks. Although we've chosen a somewhat arbitrary size for our sensor systems, exploring microfabrication technology's limitations is our fundamental goal. Because of its discreet size, substantial functionality, connectivity, and anticipated low cost, Smart Dust will facilitate innovative methods of interacting with the environment, providing more information from more places less intrusively. We use Smart Dust to pursue projects such as

- monitoring environmental conditions that affect crops and livestock;
- building virtual keyboards;
- managing inventory control;
- monitoring product quality;
- constructing smart office spaces; and
- providing interfaces for the disabled.

SMART DUST REQUIREMENTS

Smart Dust requires both evolutionary and revolutionary advances in miniaturization, integration, and energy management. Designers can use microelectromechanical systems (MEMS) to build small sensors, optical communication components, and power supplies, whereas microelectronics provides increasing functionality in smaller areas, with lower energy consumption. Figure 1 shows the conceptual diagram of a Smart Dust mote. The power system consists of a thick-film battery, a solar cell with a charge-integrating capacitor for periods of darkness, or both. Depending on its objective, the design integrates various sensors, including light, temperature, vibration, magnetic field, acoustic, and wind shear, onto the mote. An integrated circuit provides sensor-signal processing, communication, control, data storage, and energy management. A photodiode allows optical data reception. We are presently exploring two transmission schemes: passive transmission using a corner-cube retroreflector, and active transmission using a laser diode and steerable mirrors.

The mote's minuscule size makes energy management a key component. Current battery and capacitor technology stores approximately 1 joule per cubic mm

- deploying defense networks rapidly by unmanned aerial vehicles or artillery;
- monitoring rotating-compression-blade high-cycle fatigue;
- tracking the movements of birds, small animals, and insects;




Figure 1. Conceptual diagram showing the a Smart Dust mote's major components: a power system, sensors, and an integrated circuit.

and 10 millijoules per cubic mm, respectively, whereas solar cells provide 1 joule per day per square mm in sunlight and 1 to 10 millijoules per day per square mm indoors. Our optical receiver consumes approximately 0.1 nanojoule per bit and the transmitter uses 1 nanojoule per bit. We expect our analog-to-digital converter to require 1 nanojoule per sample and computations to consume less than 1 picojoule per instruction, in contrast to present processors such as the CoolRisc 81³ core, which uses 22 picojoules per instruction, and the StrongARM SA1100, which consumes approximately 1 nanojoule per instruction. These estimates demonstrate that for every sensor sample or transmission, we can perform about 1,000 8-bit operations, so it is advantageous to exchange extra calculations for fewer samples or transmitted bits. Furthermore, given our 1 millijoule per day of energy from indoor lighting, each second we can sample a sensor, think about the result, and transmit some data.

To determine our research baseline and quickly develop hardware for testing networking algorithms, we used commercial off-the-shelf hardware to build a series of wireless sensor nodes. We used either optical or radio-frequency communication models to produce one cubic-inch devices. Furthermore, other research groups have used these motes to develop a tiny operating system⁴ and deploy a 100-node network.

Sensors and motors

The multibillion dollar MEMS industry has been growing for several decades, with major markets in automotive pressure sensors and accelerometers, medical sensors, and process control sensors. Recent advances in technology have put many of these sensor processes on exponentially decreasing size/power/cost curves. In addition, variations of MEMS sensor technology are used to build micromotors; millions of these micromotors are used in commercially available projection display systems, such as the Texas Instruments Digital Micromirror Device. Micromotors, combined with Smart Dust, raise the interesting possibility of making synthetic insects (see the "Microrobotics" sidebar).

COMPUTING AT THE MILLIMETER SCALE

Traditional computer architecture design has focused on decreasing a given task's execution time.⁵ To accomplish this goal, engineers have improved semiconductor processing exponentially, increasing the transistors' speed while decreasing their size, thus allowing more complex architectures that use increased parallelism on a single die. In contrast, computing in an autonomous cubic-millimeter package must focus on minimizing a given task's energy consumption. Smaller, faster transistors have reduced parasitic capacitance, thereby

resulting in diminished dynamic power consumption. Constant electric-field scaling has reduced supply voltages, producing dramatic power reductions for both high-performance and low-energy computing because dynamic power has a quadratic dependence on supply voltage. However, constant electric-field scaling also calls for a reduction in the threshold voltage. This will result in larger leakage currents, which are already a concern in the high-performance processors to be released in 2001 that will leak amps of current. Therefore, process engineers need to keep leakage currents low, which will also benefit low-energy designers. In millimeter-scale computing, the shrinking transistor's size lets designers compact significant computing power into this small area. For example, the Intel 8088 core, originally fabricated in a 3-micron process, would only require 0.12 square millimeter after shrinking lithographically into a current 0.18-micron process, with a corresponding 100× decrease in energy/instruction.

Low-energy computation

Besides advanced microfabrication technology processes, using other techniques at every level

achieves low-energy computation. First, because we use a high-performance process but operate at low speeds, we can drop the supply voltage to the minimum level at which the devices still function; theoretically this is 0.1 volt,⁶ but for 0.5- to 0.2-micron processes it is more realistically 0.2 to 0.3 volt. To minimize current leakage, which can cause significant power consumption at the low clock rates and duty cycles that these low-energy architectures use, we can increase the channel-to-source junction's reverse bias, thus increasing the threshold voltage. Initially, adding two extra supply voltages in this package may seem onerous; however, if the mote scavenges solar power, placing two small photodiodes on the integrated circuit provides the few atto-amps per device necessary to bias these junctions. Various low-power layout, circuit, and logic level techniques have been published.⁷ Figure 2 shows a consequence of using these techniques—the worst-case energy consumption of an 8-bit adder in a 0.25-micron process.

The Smart Dust mote's tasks closely relate to the physical realm, where the fastest sampling is 10 to 20 kHz for vibration and acoustic sensors so the amount

Microrobotics

Richard Yeh, University of California, Berkeley

Add legs or wings to Smart Dust and you get microrobots. Like Smart Dust, these synthetic insects will sense, think, and communicate. In addition, they will have the ability to move about and interact physically with their environment. We can use micromachining to build microactuators and micromechanisms, forming legs and wings, which are integrated with other Smart Dust components.

The crawling microrobot shown in Figure A consumes only tens of microwatts of power; the motors can lift more than 130 times the robot's own weight. Figure B shows the flying microrobot, based on the blowfly *Calliphora*, which will have a 10- to 25-mm wingspan and will sustain autonomous flight. Developers folded 50-micron-thick stainless steel into the desired shape to create the wings and exoskeleton. Piezoelectric motors attached to the exoskeleton actuate the wings. These legged and winged microrobots will consume a total power of less than 10 milliwatts, provided by onboard solar cells. For more informa-




Figure A. Model of a crawling microrobot developed by University of California researchers. This device measures less than one cubic millimeter. Developers folded 2-micron-thick silicon sheets to create insect-like legs with micro-hinges on the folds and joints. The hollow structure has lightweight, rigid legs. Silicon tendons inside the legs couple each rigid leg segment to electrostatic motors on the robot's body.

tion, visit <http://www-sac.eecs.berkeley.edu/~yeh/currentbot.html> for the crawling microrobot project, and <http://robotics.eecs.berkeley.edu/~ronf/mfi.html> for the flying microrobot project.




Figure B. A flying microrobot model capable of autonomous flight. To create the wings and exoskeleton, the developers folded 50-micron-thick stainless steel into the desired shapes. Piezoelectric motors attached to the exoskeleton consume less than 10 milliwatts.

Richard Yeh is a PhD candidate at the Berkeley Sensor and Actuator Center, University of California, Berkeley. His research interests include microrobotics, microactuators, and microfabrication technology. He received an MS in electrical engineering from the University of California, Los Angeles. Contact him at yeh@eecs.berkeley.edu.

of data is small enough that we can use low data transmission rates. Therefore, we can use clock rates in the 1- to 100-kHz range to decrease dynamic power consumption. Despite these low clock rates, the circuits perform all their transitions during a small portion of the cycle; then they remain idle. Thus, powering down blocks for even a few clock cycles saves energy.

Remote programmability

An autonomous cubic-millimeter platform's computing requirements depend on the target application because dedicated hardware solutions usually consume less energy than a software solution. To prevent extraneous power consumption, we need to determine the minimum amount of programmability necessary for a useful platform.

The basic mote periodically samples one or more sensors, stores the values in memory, listens to an incoming packet, and transmits current or stored data. Because transmitting the data and sampling the sensors consume more energy than performing a computation, we can add more computation—such as thresholding, filtering, spectral analysis,⁸ classification,⁹ Doppler shift determination, and encryption to improve memory use and determine the significance of readings—thus providing smarter sampling rates and reducing the data transmission volume.

Remote programmability plays an important role in millimeter-scale computing. Given their small size and large numbers, we prefer to program these devices en masse, without direct connections. Remote programmability also avoids the costs of recollecting and reprogramming devices after we deploy them.

COMMUNICATING FROM A GRAIN OF SAND

Smart Dust's full potential can only be attained when the sensor nodes communicate with one another or with a central base station. Wireless communication facilitates simultaneous data collection from thousands of sensors. There are several options for communicating to and from a cubic-millimeter computer. Radio-frequency and optical communications each have their strengths and weaknesses.

Radio-frequency communication is well understood, but currently requires minimum power levels in the multiple milliwatt range due to analog mixers, filters, and oscillators. If whisker-thin antennas of centimeter length can be accepted as a part of a dust mote, then reasonably efficient antennas can be made for radio-frequency communication. While the smallest complete radios are still on the order of a few hundred cubic millimeters, there is active work in academia and industry to produce cubic-millimeter radios.

Semiconductor lasers and diode receivers are intrinsically small, and the corresponding transmission and detection circuitry for on/off keyed optical communi-




Figure 2. Hspice simulated energy consumption for the worst-case addition in an 8-bit ripple-carry adder implemented in a 0.25-micron complementary metal oxide semiconductor (CMOS) process, demonstrating the effects of supply voltage (V_{dd}) and channel-to-source junction reverse bias (V_{sb}) scaling and low-power circuit techniques. All simulation runs were performed over the same time interval and include leakage current consumption.

cation is more amenable to low-power operation than most radio schema. Perhaps most important, optical power can be collimated in tight beams even from small apertures. Diffraction enforces a fundamental limit on the divergence of a beam, whether it comes from an antenna or a lens. Laser pointers are cheap examples of milliradian collimation from a millimeter aperture. To get similar collimation for a 1-GHz radio-frequency signal would require an antenna 100 meters across, due to the difference in wavelength of the two transmissions. As a result, optical transmitters of millimeter size can get antenna gains of one million or more, while similarly sized radio-frequency antennas are doomed by physics to be mostly isotropic.

Collimated optical communication has two major drawbacks. Line of sight is required for all but the shortest distances, and narrow beams imply the need for accurate pointing. Of these, the pointing accuracy can be solved by MEMS technology and clever algorithms, but an optical transmitter under a leaf or in a shirt pocket is of little use to anyone. We have chosen to explore optical communication in some depth due to the potential for extreme low-power communication.

OPTICAL COMMUNICATIONS

We have explored two approaches to optical communications: passive reflective systems and active-steered laser systems. In a passive communication system, the dust mote does not require an onboard light source. Instead, a special configuration of mirrors can either reflect or not reflect light to a remote source; this procedure resembles how a heliograph operator bounces sunlight off a mirror to flash a Morse code message to ships—an idea traced to the fifth century BC, when the Greeks used reflected sunlight as a beacon signal. Figure 3 shows the corner-cube retroreflector (CCR)¹⁰ used to adapt this idea to




Figure 3. Autonomous bidirectional communication mote with a MEMS optics chip containing a corner-cube retroreflector on the large die, a CMOS application-specific integrated circuit (ASIC) for control on the 300 × 360 micron die, and a hearing aid battery for power. The total volume is 63 mm³.

Smart Dust. Designers have used this device, but on a macroscale, for years in laser range finding applications. A similar device helped scientists determine the moon's distance from Earth.

Passive reflective systems

In its simplest passive configuration, the passive-reflective device consists of three mutually orthogonal mirrors. Light enters the CCR, bounces off each of the three mirrors, and is reflected back parallel to the direction it entered. In the MEMS version, the device has one mirror mounted on a spring at an angle slightly askew from perpendicularity to the other mirrors.

In this position, because the light entering the CCR does not return along the same entry path, little light returns to the source—a digital 0. Applying voltage between this mirror and an electrode beneath it causes the mirror to shift to a position perpendicular to other mirrors, thus causing the light entering the CCR to return to its source—a digital 1. The mirror's low mass allows the CCR to switch between these two states up to a thousand times per second, using less than a nanojoule per 0→1 transition. A 1→0 transition, on the other hand, is practically free because dumping the charge stored on the electrode to the ground requires almost no energy.

Our latest Smart Dust device is a 63-mm³ autonomous bidirectional communication mote that receives an optical signal, generates a pseudorandom sequence based on this signal to emulate sensor data, and then optically transmits the result. The system contains a micromachined corner-cube reflector, a 0.078-mm³ CMOS chip that draws 17 microwatts, and a hearing aid battery. In addition to a battery-based operation, we have also powered the device

using a 2-mm² solar cell. This mote demonstrates Smart Dust's essential concepts, such as optical data transmission, data processing, energy management, miniaturization, and system integration.

A passive communication system suffers several limitations. Unable to communicate with each other, motes rely on a central station equipped with a light source to send and receive data from other motes. If a given mote does have a clear line of sight to the central station, that mote will be isolated from the network. Also, because the CCR reflects only a small fraction of the light emitted from the base station, this system's range cannot easily extend beyond 1 kilometer. To circumvent these limitations, dust motes must be active and have their own onboard light source.

Active-steered laser systems

For mote-to-mote communication, an active-steered laser communication system uses an onboard light source to send a tightly collimated light beam toward an intended receiver. Steered laser communication has the advantage of high power density; for example, a 1-milliwatt laser radiating into 1 milliradian (3.4 arc-seconds) has a density of approximately 318 kilowatts per steradian (there are 4π steradians in a sphere), as opposed to a 100-watt lightbulb that radiates 8 watts per steradian isotropically. A Smart Dust mote's emitted beam would have a divergence of approximately 1 milliradian, permitting communication over enormous distances using milliwatts of power.

Forming ad hoc multihop networks is the most exciting application of mote-to-mote communication. Multihop networks present significant challenges to current network algorithms—routing software must not only optimize each packet's latency but also consider both the transmitter's and receiver's energy reserves. Each mote must carefully weigh the needs to sense, compute, communicate, and evaluate its energy reserve status before allocating precious nanojoules of energy to turn on its transmitter or receiver. Because these motes spend most of their time sleeping, with their receivers turned off, scheduling a common awake time across the network is difficult. If motes don't wake up in a synchronized manner, a highly dynamic network topology and large packet latency result. Using burst-mode communication, in which the laser operates at up to several tens of megabits per second for a few milliseconds, provides the most energy-efficient way to schedule this network. This procedure minimizes the mote's duty cycle and better utilizes its energy reserves.

The steered agile laser transmitter consists of a semiconductor diode laser coupled with a collimating lens and MEMS beam-steering optics based on a two degree-of-freedom silicon micromirror, as Figures 4 and 5 show. This system integrates all optical components into an active 8-mm³ volume.

LISTENING TO A DUST FIELD

Many Smart Dust applications rely on direct optical communication from an entire field of dust motes to one or more base stations. These base stations must therefore be able to receive a volume of simultaneous optical transmissions. Further, communication must be possible outdoors in bright sunlight which has an intensity of approximately 1 kilowatt per square meter, although the dust motes each transmit information with a few milliwatts of power. Using a narrow-band optical filter to eliminate all sunlight except the portion near the light frequency used for communication can partially solve this second problem, but the ambient optical power often remains much stronger than the received signal power.

Advantages of imaging receivers

As with the transmitter, the short wavelength of optical transmissions compared with radio frequency overcomes both challenges. Light from a large field of view field can be focused into an image, as in our eyes or in a camera. Imaging receivers utilize this to analyze different portions of the image separately to process simultaneous transmissions from different angles. This method of distinguishing transmissions based on their originating location is referred to as space division multiple access (SDMA). In contrast, most radio-frequency antennas receive all incident radio power in a single signal, which requires using additional tactics, such as frequency tuning or code division multiple access (CDMA), to separate simultaneous transmissions.

Imaging receivers also offer the advantage of dramatically decreasing the ratio of ambient optical power to received signal power.¹² Ideally, the imaging receiver will focus all of the received power from a single transmission onto a single photodetector. If the receiver has an $n \times n$ array of pixels, then the ambient light that each pixel receives is reduced by a factor n^2 compared with a nonimaging receiver. Typically, using a value for n between 8 and 32 makes the ambient light power negligible compared with the electronic noise in the analog electronics.

Video camera. A video camera is a straightforward implementation of an imaging receiver. If each member in a colony of Smart Dust motes flashes its own signal at a rate of a few bits per second, then each transmitter will appear in the video stream at a different location in the image. We have implemented such a system using a laptop with a frame grabber that processes a real-time video signal in software.

We tested this system to transmit weather information from Twin Peaks in San Francisco to a video camera in Berkeley, 21.4 kilometers across San Francisco Bay. The transmitter consists of a one-cubic-inch Smart Dust mote mock-up that modulates an ordinary red laser pointer at a few bits per second,




Figure 4. Conceptual diagram of steered agile laser transmitter (side view). A laser emits an infrared beam that is collimated with a lens. The lens directs the narrow laser beam onto a beamsteering mirror, aiming the beam toward the intended receiver.




Figure 5. Scanning electron micrograph of first-generation steered agile laser transmitter. The microscope's chip combines a laser diode and ball lens with a micromachined two-degree-of-freedom beamsteering mirror. The optical path runs from the top of the laser diode's front facet, through the ball lens, reflects off the left-hand mirror plate, then finally reflects off the substrate before leaving the chip.

using only a 3.5-milliwatts of peak optical transmission power in a 2-milliradian cone.¹¹ Smart Dust's improved system will allow similar link distances, with the added advantage of automated transmitter-receiver alignment.

Using a high-speed camera and a dedicated digital signal processor to process the video signal achieves higher data rates. With modern cameras and DSPs, processing video at about 1,000 frames per second should be feasible. This would allow communication at a few




Figure 6. Pictorial representation of an integrated imaging receiver in action with predicted specifications. The receiver leverages the power of integrated circuits and CMOS imaging sensors to create a microchip with a complete asynchronous receiver circuit integrated into every pixel in the imaging array. Each pixel autonomously monitors its own signal as it searches for a transmission, decodes it, and alerts the network when it receives a data packet.




Figure 7. Typical components within each pixel of a simple integrated imaging receiver. Each of the imager's pixels contains a photosensor and circuits to perform analog signal processing and amplification, analog-to-digital conversion, and an asynchronous serial receiver. Cyclic redundancy check (CRC) is a data coding algorithm that allows a receiver to detect whether a received data packet is valid or corrupt.

hundred bits per second, which is acceptable for many applications. An alternative receiver architecture provides a more elegant solution at much higher data rates, avoiding the need for computationally intensive video processing and very high speed cameras.

"Smart pixel" integrated imaging receivers. We are developing a fully integrated CMOS imaging device that receives data at up to a few megabits per second. As Figure 6 shows, this receiver leverages the power of shrinking integrated circuits and recent developments in CMOS "smart pixel" sensors to create a microchip similar to a digital camera sensor, but with a complete asynchronous receiver circuit integrated into every pixel in the imaging array.

During the receiver's operation, each pixel autonomously monitors its own signal, looks for a transmission, decodes it, and transmits the data off chip when it receives a valid data packet. As Figure 7 shows, to function in this way, each pixel in the imager requires a photosensor and circuits to perform analog signal processing and amplification, analog-to-digital conversion, and an asynchronous serial receiver. Such a receiver should be able to receive transmissions of only a few milliwatts in strength at up to a few megabytes per second over a distance of several kilometers. Using larger collection lenses that measure more than 10 cm and high-resolution arrays— 64×64 or greater—to divide unwanted background light permits communication with low-earth-orbit satellites.

We currently have a working pixel front-end that consumes 10 picojoules per bit per pixel. It can detect 100 femtojoules per bit of optical power, but that will soon be down to 10 femtojoules per bit. This translates into the need for just 10 picojoules per bit of transmit power for short-range, indoor communication.

Integrating an imaging receiver onto a single microchip imposes severe constraints in silicon area and power consumption per pixel. Only recently have

continuing reductions in transistor size allowed for sufficient reductions in circuit area and power consumption to achieve this level of integration.

Research in the wireless sensor network area is growing rapidly in both academia and industry. Most major universities and many companies now have sensor networking projects, and some products are appearing on the market. Innovative research includes short-range micropower radio, energy scavenging from thermal gradients and vibration, operating systems, networking and signal processing algorithms, and applications. While the raw power of future computing environments will enable more massive and amazing hardware and software networks, a growing community will be pushing the limits on the lower end, building smaller hardware and writing terser code. We will program the walls and the furniture, and some day even the insects and the dust. ★

Acknowledgments

Our work is funded by the Defense Advanced Research Projects Agency's Microsystems Technology Office, the Howard Hughes Doctoral Fellowship Program, and the Fannie and John Hertz Foundation.

References

1. M. Weiser and J.S. Brown, "The Coming Age of Calm Technology," <http://www.ubiq.com/hypertext/weiser/acmfuture2endnote.htm>, 1996.
2. J. Kahn, R.H. Katz, K. Pister, "Emerging Challenges: Mobile Networking for 'Smart Dust,'" *J. Comm. Networks*, Sept. 2000, pp. 188-196.
3. C. Piguet et al., "Low-Power Design of 8-b Embedded CoolRisc Micro-controller Cores," *IEEE J. Solid-State Circuits*, July 1997, pp. 1067-1078.
4. J. Hill et al., "System Architecture Directions for Networked Sensors," *Proc. 9th Int'l Conf. Architectural Support for Programming Languages and Operating Systems (ASPLOS-IX)*, ACM Press, New York, 2000, pp. 93-104.
5. J.L. Hennessy, D.A. Patterson, *Computer Architecture: A Quantitative Approach*, 2nd ed., Morgan Kaufmann, Menlo Park, Calif., 1996.
6. J.D. Meindl, "Low Power Microelectronics: Retrospect and Prospect." *Proc IEEE*, Apr. 1995, pp. 619-635.
7. C. Piguet, "Low-Power and Low-Voltage CMOS Digital Design," *Microelectronic Eng.*, Elsevier Science BV, The Netherlands, 1997, pp. 179-208.
8. G. Asada et al., "Wireless Integrated Network Sensors: Low Power Systems on a Chip," *Proc. 24th European Solid-State Circuits Conf., (ESSCIRC 98)*, Editions Frontieres, Paris, 1998, pp. 9-16.
9. R. Amirtharajah et al., "A Micropower Programmable DSP Powered Using a MEMS-Based Vibration-to-Electric Energy Converter," *2000 IEEE Int'l Solid-State Circuits Conf.*, IEEE Press, Piscataway, N.J., 2000, pp. 362-363, 469.
10. P.B. Chu et al., "Optical Communication Using Micro Corner Cube Reflectors," *Proc. 10th IEEE Int'l Workshop on Micro Electro Mechanical Systems*, IEEE Press, Piscataway, N.J., 1997, pp. 350-355.
11. M. Last, K.S.J. Pister, "2-DOF Actuated Micromirror Designed for Large DC Deflection," *Proc. 3rd Int'l Conf. Micro Opto Electro Mechanical Systems 99*, IEEE Press, Piscataway, N.J., 1999, pp. 239-243.
12. J.M. Kahn et al., "Imaging Diversity Receivers for High-Speed Infrared Wireless Communication," *IEEE Comm.*, Dec. 1998, pp. 88-94.

Brett Warneke is a PhD candidate at the Berkeley Sensor and Actuator Center, University of California, Berkeley. His research interests include ultra-low power integrated circuits, CMOS micromachining, distributed sensor networks, and RF MEMS. He received an MS in electrical engineering from the University of California, Los Angeles. Contact him at warneke@eecs.berkeley.edu.

Matt Last is a graduate student researcher at the Berkeley Sensor and Actuator Center, University of California, Berkeley, California. His research interests include Optical MEMS. He has an MS in electrical engineering from the University of California, Berkeley. Contact him at mattlast@eecs.berkeley.edu.

Brian Liebowitz is a graduate student researcher at the Berkeley Sensor and Actuator Center, University, Berkeley. His research interests include low-power analog integrated circuits, optical communication, and MEMS. He received a BS in electrical engineering from Columbia University. Contact him at liebowitz@ieee.org.

Kristofer S.J. Pister is an associate professor of electrical engineering and computer sciences at the Berkeley Sensor and Actuator Center, University of California, Berkeley, California. His research interests include the development and use of standard MEMS fabrication technologies, microrobotics, and CAD for MEMS. He received a PhD in electrical engineering from the University of California, Berkeley. Contact him at pister@eecs.berkeley.edu.