Technical Report

Frontiers in sensors and sensing systems

Steven D. Glaser†

Center for Information Technology Research in the Interest of Society, University of California, Berkeley, U.S.A

Rahmat A. Shoureshi‡

School of Engineering and Computer Science, University of Denver, Denver, CO, U.S.A

David Pescovitz‡†

Science Writer, San Francisco, CA, U.S.A

1. Introduction

Recent technological advancements in materials science, micro fabrication of MEMS (microelectromechanical systems), and bioengineered systems have made the dream of inexpensive, powerful, ubiquitous sensing a readily achievable reality. Examples range from truly smart airframes and self-evaluating buildings and infrastructure for natural hazard mitigation to large-scale weather forecasting and self-organizing energy systems. A common thread running through these, and all other applications of ubiquitous sensing, is the vast amounts of data generated, and a need to have the ubiquitous sensor networks process this data in order to return decisions and information. Sensor networks become a dynamic organism far more powerful and user-friendly than the traditional view of a sensor as a widget, an individual component that needs to be deployed, programmed, and interrogated. Therefore, the convergence of sensor technologies, communications, and computing has the potential to overcome barriers of time, scale, materials and environment.

The National Workshop on Future Sensing Systems was held in Lake Tahoe, California, August 26-28, 2002, sponsored by the Sensors Technology Program of the NSF Division of Civil and Mechanical Systems (CMS), with co-sponsorship from DARPA, NIH, DOE, NIST, NASA, AFOSR, ONR, ARO, ARL, and NRL. The scope of this workshop encompassed discussions of research needs, current and emerging technologies, and efficacious partnerships required to develop and implement future sensing systems. The workshop was planned to stimulate synergistic ideas and directions generated by industrial, scientific and government participants, and to start planning a long-term road

†Professor, Corresponding Author, E-mail; glaser@ce.berkeley.edu
‡Professor and Dean, E-mail: rshoures@du.edu
‡†CEO, E-mail: david@pesco.net

map for R&D projects related to emerging needs and technologies. Furthermore, this workshop started strategic partnerships among industry, scientific community, and government agencies that are developing innovative and cross-pollinating sensor systems capable of converting raw data into useful information about myriad environments.

The workshop was designed to push the community to share and work together, not study the current state of the art. This gathering was not a place for the participants to give formal presentations about their research successes. The participants drove this workshop. The ideas, the concepts, and the research directions outlined in this paper were decided by all the participants.

The workshop spawned numerous collaborations, proposals, and plans for future sensor research. It also proved that there are far more commonalities than differences in the challenges faced by a wide variety of researchers from the sensor community at large. Perhaps most importantly though, the workshop helped foster the strategic and targeted growth of that nascent community.

The manuscript will first discuss the Workshop and its goals and organization, followed by an exposition of the research needs and directions of future growth in sensor-directed research. The final section discusses some proposed changes to how the research community might better organize to maximize the research output per dollar invested, and deliver more useful results. A list of defined abbreviations used in this paper is given by Appendix 1. Please note that the full word-for-word transcripts of the Workshop, as well as all the presentations, is freely available on-line at http://www.ce.berkeley.edu/Programs/Geoengineering/sensors/.

2. The national workshop on future sensing systems

2.1. Workshop rationale

The National Workshop on Future Sensing Systems was a vehicle to start researchers on the road to integrating the possibilities in development by the many disciplines working in the common space. In the past, advances have often been duplicated amongst the different disciplines and end users due to lack of communication. An example is the desire of the structural engineering community to develop a custom sensor when commercial off the shelf (COTS) devices already exist that can fill most of researchers' desires very inexpensively. The scope of this workshop encompassed discussions of research needs, current and emerging technologies, and efficacious partnerships required to develop and implement future sensing systems. The workshop was planned to stimulate synergistic ideas and directions generated by industrial, scientific, and government participants, and to start planning a long-term road map for R&D projects related to emerging needs and technologies. Furthermore, this workshop aimed to encourage strategic partnerships among industry, scientific community, and government agencies that are developing innovative and cross-pollinating sensor systems capable of converting raw data into useful information about myriad environments.

2.2. Workshop goals

The principle deliverables of the National Workshop on Future Sensing Systems were as follows:
1. Produce a report to aid in charting research needs, opportunities, and directions for future growth of the NSF research portfolio in sensor technologies.
2. Provide, as a resource to the CMS community, information about industrial and federal agencies'

Table 1 Affiliations of the attendees

Academic discipline	
Civil engineering: 28	Electrical engineering: 26
Mechanical engineering: 26	Computer science: 9
Chemical engineering/chemistry: 9	Biomedical: 8
Employment discipline	
Government: 29	Academia: 73 (students: 15)
Research laboratories: 15	Industry: 10

   programmatic interests, scopes, budgets, etc. for basic and applied research on sensors, sensing materials, and systems-level sensor integration issues.

3. Provide recommendations and suggest mechanisms for structuring interdisciplinary and interagency coordination on joint sensor systems research initiatives.

4. Stimulate increased domestic and international cooperation among academic, industry and Federal researchers through exchange of ideas, interests, and program information, and suggest new ways to initiate collaborations.

The workshop itself was the first step towards fulfilling the fourth task. The discussion forums brought a wide variety of researchers involved in sensor research to the table, literally. During breaks, presentations, and open time, participants began to note the underlying interests they all share. They took steps to rise above the traditional, superficial differences that often keep researchers from myriad fields from collaborating as a matter of course. As shown in Table 1, the workshop attendees represented an impressive mix of disciplines, with an even mix of electrical, mechanical, and civil engineers, as well as chemists and computer scientists. Other fields included seismology, mining, aerospace, material science, physics, and engineering mechanics. Experimentalists and theoreticians were well represented. Twelve of the attendees were lured from foreign countries. There were representatives from government agencies, funding and technology consumers, multiple national laboratories, both small and large business, and academia.

There were a variety of different agendas for the attendees. The goals in organizing the workshop were to facilitate people with different backgrounds working together, and fostering examination of what they share in common. Through the formal and, equally important, informal discussions, individuals had the opportunity to think about sensor applications broadly, or focus deeply on specific technical challenges. These conversations sparked numerous ideas and possible solutions. For this reason alone the workshop was an exciting, interesting experience for everyone.

NSF had goals, chiefly to produce a report to help form the directions of the new initiatives in sensing. NSF wanted to provide a resource to help the community realize what different sources of funding, programs, and mechanisms are in place to aid research in sensors. In addition to increased domestic cooperation, the NSF particularly wanted to increase cooperation with research entities from the EU, Asia, and other countries. Finally, each participant had goals that might be very different from the NSF's goals or the organizers' goals.

2.3. Background and organization

The workshop was designed to push the community to share and work together, not necessarily to study the current state of the art. This gathering was not a place for the participants to give formal

presentations about their research successes. There are ample venues for this, and all can read the published literature.

There was a variety of deliverables for our sponsors and community, some of which were suited for a formal setting, and some not. As a result, the workshop had a variety of formats. The first day of the workshop was devoted to presentations about technologies that might lead the research for the next decade. The speakers were asked to selflessly present a little of what is possible today, and a lot of what the community might strive towards. The talks were meant to expand minds, and be a learning experience for everyone (the speakers included), so that the participants had a common place to begin discussions.

There were two panel/open discussion sessions: Venture Capitalist and Industrial View of Investment in Sensor Technology Research, and Formation of Research Partnerships. The workshop closed with a session entitled Agencies' Presentations on New Initiatives in Sensor Systems. Much time was spent in "small" (about 20 participants) discussion groups, to which participants were assigned so that all could benefit from the great diversity of backgrounds and knowledge to be shared. The topics of the breakout groups were:

- Data integration -- how to handle and most efficiently use the vast volumes of data that will be generated by future implementations of ubiquitous sensing
- Smart devices and self-assembling networks--how to use this new paradigm, how to comprehend previously impractical applications
- Integration of new sensing and communication technologies into practice--COTS and research
- Symbiotic synthesis of constructive modeling, interpretation, and sensing
- From science and technology to useful and affordable sensors

All groups also had to make the case that research spending for advanced technologies for civil infrastructure (in its broadest sense) is cost effective.

3. Research needs, opportunities, and directions for future growth of the research portfolio in sensor technologies

In the past, sensors have tended to be specification-driven replacements for electromechanical devices, components built for high-volume applications (e.g. pressure sensors, airbag accelerometers) that end up as commodities driven toward very low costs. Sensors are now evolving toward systems for addressing high-level, problem-driven needs and applications. Sensors are becoming solutions rather than devices. Common characteristics of sensing systems include integration of multiple sensors, multiple system concepts/approaches, autonomous operation, flexibility, onboard intelligence, and the growing use of wireless interfaces. A cartoon illustration of this concept is shown in Fig. 1.

The present push in sensors and sensor systems is to make extensive use of MEMS-based devices. MEMS, or micro-electromechanical systems, are micron-scale sensing devices built using the etching technologies devised for manufacturing integrated circuits. The promise of MEMS, circa 1980, is co-located sensing, actuation, measurement, signal processing, computation, and communication based on the IC manufacturing paradigm. The reality, circa 2004, is very limited progress toward either systems-on-a chip or systemsin-a-package. Technical barriers include the fact that co-fabrication is presently difficult, making on-chip solutions expensive, and that packaging technologies are just emerging for sensors other than inertial and optical devices. Economic issues include lack of truly high-volume applications that would establish the infrastructure, and a lack of industry-wide standards.
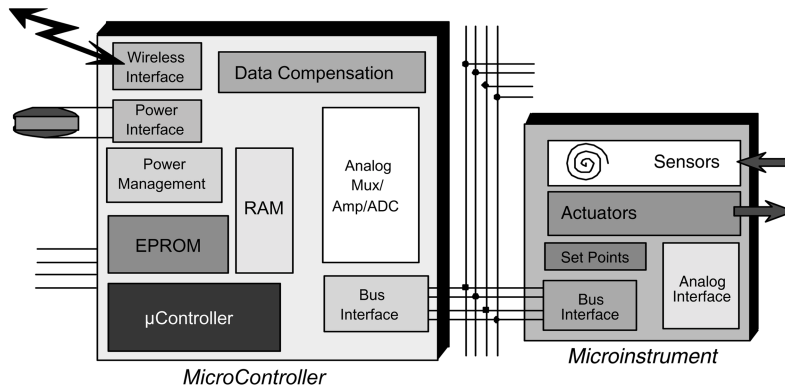
Fig. 1 Example of a sensor system: Environmental Monitoring System, senses many parameters, e.g. temperature, humidity, airspeed/direction, gases, chemical vapors, etc.

Co-fabricating electronics and microstructures is not easy. For example, mixed MEMS and complimentary metal oxide circuitry (CMOS) is a "boutique process", making it quite expensive. The "MEMS first, CMOS last" approach requires that you have, or control, your own CMOS fabrication facility, e.g. Sandia and Analog Devices. The "CMOS first, MEMS last" approach might be the best approach, but high thermal requirements for MEMS fabrication makes it difficult not to cook the CMOS in the process.

The automotive, aerospace, petroleum, chemical, biological, nuclear, military, and civil users need devices that can operate under high temperatures subject to intense vibrations and corrosive and abrasive media, and perhaps high radiation. Silicon, the basic building block of MEMS, is not desirable for harsh environments because electrical properties degrade above 350ºC and mechanical properties degrade above 650ºC. Silicon is also susceptible to radiation damage and etched by corrosive chemicals.

However, the rate of advancement in MEMS sensors is profound, with thousands of new commercial devices coming to market every month. There are many advances in networking IT, driven in part by DARPA (i.e. the NEST project). Given the present and future possibilities, and the very low cost offered by mass production, MEMS-based sensors, particularly with networking capabilities, are the future. These are the devices the organizers assumed to be used by all the workshop attendees, and therefore are the focus within this report.

3.1. Cost

In order for the benefits of ubiquitous sensing to be truly realized, the sensors themselves must be inexpensive enough for researchers to actually deploy them in great numbers. This is true whether the intended applications are smart airframes, office buildings that self-evaluate structural integrity for natural hazard mitigation, or animal habitats where the ecosystem must be carefully monitored to provide environmental insights. "Usefulness" and "affordability" must be judged in the context of the needs of applications and users. Needs for new useful and affordable sensors are being driven in part by the development of the additional technologies required for low-cost distributed sensor networks, specifically wireless communications, low power electronics, and power sources. To meet the needs of low-cost distributed sensing networks, sensors may need to be smaller, lower-power, simpler, and lower-cost than existing sensor technologies can provide.

While bulk production would certainly lower the cost of sensors, the key to appreciably reducing the price is merging MEMS and electronics onto the same substrate. Following the path of integrated circuit manufacturing, batch fabrication emerged in the 1980s. The reality though is that while batch fabrication methods have rapidly improved, systems-on-a-chip are still in their infancy, limited to so-called boutique processes where each step is married together for a specific chip and a specific application. The small number of manufactured pieces result in high cost per piece. Following the natural analog of the IC industry, only devices with a large perceived market will become commercially available at a very low price. Researchers in all disciplines should therefore closely examine their options in commercially-available off-the-shelf devices to expand the economy of scale.

According to several workshop participants, a modular process for sensor fabrication is one way to keep costs down. As briefly mentioned in the previous section of this report, many researchers are investigating fabricating MEMS first followed by CMOS, traditional integrated circuit technology. Others take the opposite approach, integrating gyroscopes or optical MEMS onto CMOS devices. Materials like silicon germanium alloys are ideal for this purpose as they provide strong bonding, high quality, and can be deposited on the CMOS at much lower temperatures. Packaging the sensors to protect them from the potentially harsh environments where they will be deployed presents another fabrication concern. The packaging must protect the electronics at the same time allowing the sensors to take unobstructed measurements. The sensors themselves must also be capable of some level of automatic calibration in the field where dust or other substances can coat the sensing elements. While often neglected, packaging could become as large of a challenge as device fabrication itself. This is certainly true for the IC industry. Doing a proper job of packaging requires a systems approach, with vertical integration of all levels of designers and users.

Another aspect of cost is the price of system integration, costs over and above the price of a sensor chip. The first step to reducing the price is the transformation from sensor to instrument, the integration and processing of several sensors at a given location. It was pointed out that building systems involves a certain amount of work that might not be considered part of "basic research" for a given discipline. This attitude needs to change. Here again, vertical integration of the various skills needed to construct and operate a sensing system has to happen from the project inauguration. Honda's team-based effort to construct its highly instrumented auto-adaptive engine is a quintessential success story of good vertical integration of skills.

In the case of sensors, further research is necessary to identify the best design and related materials and fabrication approaches suited for parallel assembly and encapsulation in a foundry at an affordable cost. These challenges demand collaboration between electrical engineers, mechanical engineers, materials scientists, and, of course, representatives from the myriad disciplines who ultimately will use the devices. Indeed, cutting costs is a multidimensional problem that must be solved if sensors are to be distributed en masse for the applications discussed by the workshop participants.

3.2. Ease-of-use

Not only must the sensors be inexpensive, they must be easy to configure and deploy. This problem is two-fold. The first challenge is developing hardware that can easily be configured to suit the needs of the application. Modularity of the sensing components is one possibility. While entire systems-on-a-chip can reduce cost, niche applications will require a very specific configuration of sensors. To that end, a communication/computation platform or separate platform segments could be manufactured in bulk. Then, a selection of sensors could be connected to the communication/computation platform to

satisfy the sensing requirements of the particular application. This would not necessarily result in the most seamless, elegant, "perfect" solution. However, this approach would be the most empowering to the end users, freeing them from sole-source hardware solutions.

If the systems are to be deployed in remote environments--uninhabited nature preserves or inside a buildings walls, for example--the networks themselves must be self-configuring. Firstly, if the sensors are to be deployed in the thousands, manual-configuration is simply not practical. Techniques are being considered that "virtually" configure and program the sensors with new software, injecting new code into the network and letting it proliferate from sensor to sensor. The extreme flexibility offered by sensor systems based largely on reconfigurable software is a double-edged sword. There is a near-infinite variety of detailed applications that can be addressed and number of possible system reactions to events. The power of this variety is matched by the difficulty in low-level programming of such sensing devices. The learning curve has proved to be very steep for users interested in real applications. In addition, the software/ operating systems are still in their infancy, and there is significant difficulty in developing stable and reliable applications for any user. Perhaps one immediate solution is for the software developers to write simplified toolboxes of commands commonly used for instrumentation and experimentation. In actuality, a set of commands for this purpose would be quite compact because there is great commonality among expected experimental users.

Additionally, the architectures of these future sensing networks are quite different than familiar network architectures like that of the Internet. The real world is noisy and unpredictable. Tight power constraints may require that the sensors radio is almost always off, switching on only to transmit or receive data in short bursts. This means that the network itself cannot transmit much information about its own configuration or health. What is left is a probabilistic connectivity model. As a result of only having partial information about the network as a whole and the individual sensors themselves, new network algorithms are required.

3.3. Applications

Applications of sensing span nearly every engineering and scientific discipline. Civil engineers are particularly interested in instrumenting structures to monitor stability after earthquakes and other extreme events, and also as tools for structural health prognostication for the built infrastructure. If the sensors are inexpensive and robust enough, they can literally be installed inside the walls of a structure to gather data about how seismic waves move up a building. Some of the application areas highlighted by the Workshop are listed in Table 2. The uniqueness of civil structure systems, e.g. a bridge, is their life span of over 100 years. Thus, developing embedded sensors and networks that can operate for a century is a daunting challenge.

NASA researchers are interested in furthering structural health prognostication to monitor space vehicles, in particular a vehicles thermal protection system. The aim here is to electronically inspect a vehicles structure so that it can return to space within twenty-four hours of landing, far speedier than the six months it took to turn around the Space Shuttle. To do this, high temperature sensors are required. A similar effort is underway to develop a system that detects structural "hot spots" in an aircraft body that may suffer from corrosion or cracking. Embedding sensors within the structure of the aircraft could eliminate the high-cost and time necessary to de-skin a wing to inspect it.

Mark Derriso of Wright-Patterson Air Force Base presented related problems identified by the Air Force that are common to many engineering applications. Currently, satellite structures are made almost exclusively of composites, primarily carbon fiber-reinforced thermosetting polymers. In addition to their

Table 2 Some examples of application areas

Health monitoring of civil infrastructure, bridges, buildings, pipelines
Effects of extreme events in infrastructure
Security monitoring of critical infrastructure
Distributed environmental monitoring systems
Search and rescue systems
Process control
Biomedical monitoring
Homeland security
Supply-chain management
Smart buildings
Inventory tracking
Traffic monitoring/control
Energy-use monitoring/control
Aircraft/spacecraft structural health and control

high specific stiffness and strength, carbon fiber composites have the great advantage of possessing extremely low coefficients of thermal expansion. Thus, they are used in antenna structures, sensor supports, and optical benches. It is well-known that residual stresses inherent in the laminate due to cool-down from the curing temperature, as well as the stresses caused by thermal cycling in orbit, can initiate and propagate microcracks which results in dimensional instability. The availability of sensors that can be embedded would allow for the monitoring of the initiation and growth of microcracks over the life of the spacecraft. By using a sensor array, the location and size of the micro-crack can be ascertained and, using mechanics of laminated materials (Lamb wave theory), the effect of the micro-cracks on distortion of the structure can be determined. Such information is invaluable for determining the effect of such behavior on critical satellite operational parameters such as pointing accuracy. These issues are applicable to any fiber-reinforced composite structure, on Earth, in the atmosphere, or in space.

One of the earliest non-engineering disciplines at the forefront of sensing technology applications was biology. A large variety of chemical and biological sensors have been developed. Many of these devices are directly relevant to a large part of the engineering community, e.g., geoenvironmental, but have largely gone ignored. This is an area where improved communication and joint research efforts could have tremendous benefit. Sensors are currently being tested in tree canopies to monitor the micro climates at varying heights in forests. Meanwhile, Intel Research is involved in a collaboration to deploy hundreds of sensor nodes on an uninhabited island off the cost of Maine to monitor an elusive species of nesting seabirds. Similar sensor and data manipulation technologies can aid environmental scientists in understanding how contaminants travel through soil.

As sensing technology is improved, myriad other applications arise. Portable chemical and biological sensor arrays are in development to detect chemical and biological weapons "on the spot" in airports and other high-risk facilities. The Jet Propulsion Laboratory is calling for sensors to be integrated into drill bits to analyze rocks on Earth and other planet, perhaps even to detect biological markers that signify the presence of extraterrestrial life. Instrumented drill bits could also aid crime investigators in determining if bodies are buried in concrete or walls. In addition to instrumenting

equipment with sensors to expand its functionality, sensors would also be useful for equipment monitoring itself. For example, outfitting high-value machinery, from HVAC chillers to jet engines, could aid technicians in diagnosing the state of the system and predicting the expected longevity of its components.

3.4. System design

In order to develop useful sensors, it is critical to first understand the physics of the parameters to be measured. Only then can appropriate sensors, with the necessary operability, be developed. In fact, many sensing problems and needs are chemical in nature, and require understanding of the chemistry of the parameters to be measured. Physics and chemistry are also important for the development of other components of distributed wireless sensing systems and networks, such as power sources, and radio frequency (RF) or optical communications.

The arena of wireless (RF) sensor nodes generated the most technical interest and creative thought at the workshop. A generic wireless sensor platform contains several key components, each with its own set of unique design challenges: an antenna, a power supply, a transceiver, signal processing, and a microprocessor running the sensing and networking software. The entire device must be robust, cost-effective, and potentially quite small. All of these are major sensor design challenges. There are many research opportunities in the area of RF MEMS, ranging from high Q resonators and low-loss switches to tunable components enabling the bandwidth and center frequency of the radio to be changed in situ.

It must be noted that wireless is not always the only or best solution for a given set of problems. RF uses great amounts of power compared to MEMS-based sensors and modern microcontrollers (MCUs). There are problems insuring privacy of the data transmission that have not yet been satisfactorily addressed. The quality of low power RF transmission is quite variable and susceptible to seemingly insignificant changes in surroundings. Such variables include orientation of antenna, distance from antenna to structural elements, proximity of water-bearing bodies, fog, etc. Data transmission rates over low power RF networks are often quite limited compared to the simplest wired bus. A single receiver can only receive one signal at a time, making true real-time operation impossible. A solution is for each sensor node to have permanent buffer memories and package the data to be sent into time-stamped, error-corrected packets to be transmitted in turn. This transmission could occur in pseudo real-time if the bandwidth is sufficient for rapid downloading of multiple sequential data streams. Accurate time synchronization in the network is then required.

Another issue often ignored is that it is impossible to transmit information-carrying RF signals through earth materials. Simple substitution of the relevant properties into the Maxwell equations and applicable equations from information theory will show that a wavelength long enough to travel, for instance, 100 meters through soil or rock, will have a frequency too low to carry any information.

Numerous approaches to powering tiny sensors are being considered. Power requirements depend greatly on the application. For example, a sensor on a missile would only draw power for a very short period of time. On the other hand, a sensor network to monitor a building's structural health would ideally be self-sustaining for a decade or more. Batteries may be suitable for some applications. Solar power supplies can be effective in well-lit environments. Progress is also being made in power-scavenging--using piezoelectric techniques to convert ambient motion, the vibration of HVAC ducts for example, to generate power. Battery life and energy density are expected to improve with time. However, battery improvements have not kept pace with the rest of the electronics industry.

3.5. Data management and interpretation

Data is valuable only if it can be transformed into useful information. Given thousands of sensor devices, what will researchers do with all the data that's streaming in? Who is going to analyze them? To manage the firehouse of data, sensor network researchers might look at data integration in a broad sense, processing and concentrating the data locally within the network. In the end, the user wants information, not numbers. One approach is to synthesize our design models with the data to be collected-- sensing is a link between the real world and an abstract world. The user can then predict through simulations the expected system properties and whether the data interrogation is correct. This synthesis leads to adaptive data interrogation systems where the interpretive section of the loop can tell the sensor system to record on a different granularity or measure different variables. In any instance some aspects of the system--sensors or model--will be system independent while others will be system specific.

Data collected by sensor networks is multi-dimensional. A single sensor can gather data about temperature, humidity, light, and a host of other variables. With power and bandwidth at a premium, embedding some local processing capabilities within the sensor networks is desirable. For example, the temperature data gathered from numerous sensors could be fed into one or more other sensors on the network for processing. A weighted average could then be calculated and transmitted to the user, significantly reducing the amount of data flying around the network. Furthermore, database systems must enable researchers to query the sensor network as a user- friendly virtual database rich with raw information about the real world.

Experimental sensor network databases like TinyDB (Intel Research/UC Berkeley) enable a user to gather the information he or she desires by posing a simple query in a familiar database language. TinyDB's declarative query language allows the user to describe the desired data--the average noise level in a building, for example--without having to tell the software how to acquire that data. The data processing is embedded in the network.

At its most fundamental level, data interpretation requires knowledge of the uncertainties involved in taking the measurement. Perhaps because of the youth of the field, quantification of error and uncertainties in a sensor network has been neglected. This is a problem that needs to be addressed. Hardware solutions must include auto-calibration so that rational estimates of uncertainty can be made. System solutions must include sensor output validation with independent measurements, thorough integration of sensors, and predictive modeling.

3.6. Sensor networks

The dream of sensor networks is described well by David Culler of UC Berkeley: Spread thousands of wireless sensor nodes casually over an arbitrary area of interest. The sensors then self-organize into a network conveying arbitrary information from any point to any other at whatever bandwidth is demanded. All the while, the sensors operate at incredibly low energy usage (i.e. off most of the time), running for years on small batteries and harvested energy. They're extremely responsive in times of key activity, without ever bothering the deployer about design considerations, intended usage, faults, or constraints.

Right now, there is no ad-hoc, multi-hop, robust, low-power network that meets all these "requirements", and getting there is a difficult challenge. However, there is a host of interesting and important research that is moving sensor networks toward a more realistic version of Culler's scenario. Meanwhile, real applications can be achieved along the way, enabling the technology to evolve in response to actual,

rather than perceived demands.

Smart devices that self-assemble into networks open up a paradigm that the research community is not necessarily familiar with. The paradigm of sensor networking allows engineers and scientists to move beyond the idea of a sensor as a single little instrument that measures one thing to a system of many small nodes working cooperatively. Once the sensing problem is examined as a system, the determination of optimal physical locations becomes an important problem. Where must a sensor be placed to accurately measure what is requested? How must the physical interface between sensor and structure be constructed? The user must decide up-front whether the data will be used for increasing basic understanding of some phenomena or for monitoring any change of state.

Due to inherent and severe constraints, the networking component of future sensor networks will look very different from the power-hungry, sparsely-instrumented IT world. Sensor network constraints include noisy environments, dynamic situations, and user access to only partial information. Low power consumption requirements mean that the radio is rarely on. However, the self-organized and distributed nature of the network provides opportunity to work across layers of abstraction, and presents a rich set of open problems.

In the huge realm of sensor research, applications define the valid assumption sets for network design and guide the front of progress. Data traffic within the network can be any-to-any, all-to-one, one-to-all, or a collection of sub-paths. Data rates can be a steady low-bandwidth steam of readings/findings, bursts of data, or periodic logs. Duration, or system lifetime, can be years (hazard alarm), months (field season), or minutes. The available infrastructure will vary tremendously among applications or within a single application. Other variables include options and constraints in power, base-stations, distribution (all fixed, all mobile, or a mixture), time synchronization, localization, and proximity. For generality and flexibility, the network should be able to handle heterogeneous data and account for time delays and system latencies. A sensor network also allows a trade-off to be made between utilizing more, less accurate sensors. For instance stacking several $8 accelerometers at a single location might yield the same accuracy as a single $800 device, a savings of two orders of magnitude.

"In order to be unattended, the systems have to be self-configuring. From the perspective of a computer scientist who has been involved in internet-related design for many years, this is a very different context. You can't just take a lot of the technology that works quite well in the context of the internet and common distributed systems and assume that they'll apply here..." (D. Estrin).

Already, humidity sensors, barometric sensors, and dozens of other measurement devices can be bought off the shelf. If these devices are integrated with a receiver and microcontroller, their usefulness dramatically increases. The sum of the sensor data is much more than the data each sensor provides. Suddenly, valuable information can be pumped from the network instead of massive amounts of raw data. At this point, the network becomes the sensor.

3.7. Test beds

It was universally agreed that developing well-documented test beds should be the top priority of NSF and the other funding organizations. An effective test bed should demonstrate a practical and complete sensing network, including robust wireless communications and long- lived power sources. These test beds should be used to tackle meaningful, but solvable problems.

If the network is indeed the sensor, utmost importance must be placed on putting entire networks through their paces. While some early experiments can be conducted in a laboratory environment, most of the hurdles in building a working system can only be revealed, and repaired, by deploying it in situ.

Workshop participants focused on such diverse problems as chemical sensing, seismic stability, and environmental monitoring all stressed the necessity of field studies. A sensor dropped in a burrow to monitor the activity of nesting birds is of no use if the birds chew through its circuitry. A solar-powered sensor network in a corn field is not helpful if there is less sunlight than expected. The goal of researchers then is to identify real world applications where sensor networks may provide value and actually deploy those networks, even on a small scale. That way, the participants agreed, the hardware and software can be put through its paces.

However, it is not enough to gather data that engineers find interesting. The hardware and software researchers must collaborate with the individuals who actually will use the data in the course of their research. Again, a cross-disciplinary approach and vertical integration of skills is the only way this is possible. Sensor network researchers must team up with integrative biologists, chemical engineers, structural engineers, physicists, etc. in order to find out what kind of data they want to see from a particular network. Then, once the hardware and software kinks of the network are worked out in the field, the data must be studied to establish its accuracy as compared to data collected using other methods of observation.

A sample of test beds in development or operation include:

A sensor network of "motes" peppering a remote island off Maine to aid in the study of the nesting patterns of petrels. (Intel Research/UC Berkeley)

A 65-sensor mote grid deployed in a vineyard to guide irrigation and planting. (King Family Farms/ Intel Research/AgCanada)

Seismometers installed on the San Andreas Fault to calculate the depth of the fault, locate accumulating stress, and potentially improve earthquake prediction. (UCLA Center for Embedded Networked Sensing)

A sensor network on the historic Masada mountaintop palace in Israel to aid in determining why the massive limestone building blocks are shifting and cracking. (UC Berkeley)

Bridge monitoring wired sensor arrays. (UI-Chicago)

Monitoring of tunnel wall deformations, London Underground. (MIT)

A variety of installations to explore structural health prognostication data gathering and analysis techniques. (LANL)

4. Programmatic needs of participating governmental agencies

4.1. National science foundation

Dr. Filbert Bartoli, chair of the NSF's Sensors Working Group, presented the NSF Sensors initiative. His working group contains members from each of the six divisions within the Engineering Directorate. The broad interest in new sensor technologies across NSF has led to a multi-year solicitation in sensor and sensor network research. The intent is to address sensing and sensors on both a component and system level for data gathering, information management, decision-making, and action.

The primary goal of the NSF is to advance fundamental knowledge of sensors and sensor systems, in any application area, with some particular thrust areas listed in Table 3. The approach presented focused on three categories. The first is new sensors and sensor components, the second is systems and networks, and the third is interpretation and use of sensed data. Of these three areas, sensor networks will clearly be the immediate growth area. A prudent reader will note that this view conforms to the traditional disciplines of electrical engineering, computer science, civil, and mechanical engineering. The

Table 3 Some interest areas for NSF sensor initiatives

MEMS sensors
Sensor networks
Packaging
Data visualization
Distributed sensing
Sensor security
Information and decision systems
Very broadband wireless communications
Integration of ubiquitous computing
Intelligent sensors and data self-management
New sensor functionalities and improved performance
Integration of sensors into engineered systems
Integration of sensed data into the decision-making process

emphasis is not necessarily on establishing multidisciplinary research through specific applications.

4.2. National institute of standards and technology

The National Institute of Standards and Technology (NIST) is primarily a research organization but funds extramural research through the Advanced Technology Program (ATP). The mission of ATP is "to accelerate the development of innovative technologies for broad national benefits through partnerships with the private sector." The ATP has an annual budget in the tens of millions of dollars, with an average project receiving $1 million per year for up to three years. The investigators own the intellectual property, and ATP helps integrate the business and technical aspects of the concept. The review of the proposal is judged on a 50% technical and 50% business-plan basis. The technical aspect must show a high amount of innovation above the present state of the art, but be based on sound physics backed up by preliminary data. The business plan must show broad economic benefits on a national scale, and make a solid case that ATP is the best source of funding to develop and grow the business.

When initiated, ATP was focused on large corporations that could match the NIST grants with millions of their own dollars. In recent years there has been a change of focus to include small firms and university researchers who want to undertake commercializing very innovative technologies. The segment of ATP most relevant to the sensor community is the Smart Advanced Materials and Structural Systems. After 9-11, there has been an active effort to fund research relevant to homeland security, an area in need of new and cost-effective sensing technology.

4.3. National institutes of health

An extremely large source of funding for advanced research is the National Institutes of Health (NIH). The NIH has more than 27 funding entities and an annual budget of $27,000,000,000 (billion!) of which 85% goes to funding extramural research. The engineering community has traditionally ignored NIH funding, and, in fairness, the NIH has ignored the engineering community. An attempt to bring expertise in physics, engineering, material science, and mathematics into the biomedical arena has led to the

forming of the National Institute for Biomedical Imaging and Bioengineering (NIBIB), the goals of which are to develop fundamentally new knowledge, foster new technologies, facilitate cross-cutting research, and nurture a new generation of researchers able to reach across traditional disciplinary boundaries.

The NIH proposal process is a bit more "onerous" than NSF, since the proposed work must be presented as a hypothesis to be proved, with some direct connection to biomedical problems and realities. However, the 2003 budget for NIBIB is $271 million. NIBIB currently is interested in four main approaches: hypotheses testing, novel designs, specific problems, and new technologies. The initial thrust areas are biomaterials, biosensors, nanoscience, bioinformatics, biomedical imaging, and image guided therapies and intervention. The director of NIBIB would very much like to see proposals from the engineering sensors community. The program is new and acting director Dr. Donna Dean is refreshingly open to unique methods and goals. As Dr. Dean said several times at the workshop, "There is a lot NIBIB does not know".

4.4. Air force research laboratory

The federal government is possibly the largest owner of manufactured infrastructure, with a large percentage comprised of military assets. In particular, the Air Force has many serious performance and maintenance issues involving aircraft. Major Mark Derriso of the Air Force Research Laboratory at Wright Patterson AFB presented some current and future needs for new and advanced sensor systems. In particular, there is a need for structural health monitoring and prognosis, a need reiterated by speakers from Los Alamos, JPL, the European Community, and many other workshop participants. This research has major commonalities across all the engineering application areas and is ripe for interdisciplinary work. Examples include field testing of bonded repairs to FRP structures, and localizing zones of fracture and corrosion growth. This research is important to the transportation infrastructure community as well as to aerospace engineers and biomedical personnel. Some other areas of concern to the Air Force that are of direct interest to sensor researchers include high-temperature sensors of most kinds, high- reliability sensors, optimization of sensor systems to minimize weight, and reliable wireless technology. These problems involve common aspects and should attract the talents of sensor engineers from the public and private sector. These are real and challenging problems that must be answered with reliable, cost effective solutions.

4.5. Ministry of economy, trade and industry, japan

Dr. Yoshikazu Goto of the Japanese Ministry of Economy, Trade and Industry (government of Japan) also provided workshop participants with an interesting and informative look at METI's needs related to sensor research. In particular, METI hopes to spur interest in the maintenance of existing infrastructure and plants as opposed to building new structures. As a maturing economy, Japan must find a way to prosper by meeting society's daily needs in a sustainable way (as opposed to unbridled growth). This includes environmental restoration and preservation, public health and safety, and extending the usable lifetime of existing infrastructure. These issues are relevant to Europe and North America as well. As mentioned by virtually every speaker, sensors could play a pivotal role in Japan's restoration and preservation plans.

METI believes that the greatest funding opportunities for researchers is to be gained by realizing the overwhelming need to integrate sensing technologies with the major technological thrust areas of information technology, biotechnology, nanotechnology, and the environment. To this end, test beds

become a vital way to come to standardization of data, acquisition, and sharing, as does vertical integration of science, technology, industries, operators, and owners. Both of these areas were discussed in depth at the conference and outlined in this report.

5. Some recommendations and mechanisms for structuring interdisciplinary and interagency coordination on joint sensor systems research initiatives, and stimulating cooperation among academic, industry, and federal researchers

5.1. Vertical integration of skills

One of the conference goals was to make a case for vertical integration. There is a traditional model of funding where individuals, perhaps electrical or mechanical engineers, ask for funds to develop a single device. In the end, that sensor was developed in isolation and rarely reaches the laboratories of those who would use it for real applications. Prof. Jiri Janata talked about 30,000 papers on chemical sensors being published each year, yet very few such sensors are available to an interested user. What if these pieces were developed with some vertical input up front?

To ensure that new sensors are developed to address important needs and are useful and affordable, it is suggested that the user groups involved in important applications (academic disciplines, national labs, etc.) be involved both in writing requests for proposals for new sensing system developments, and in evaluation of the proposals. It is desirable to get understanding and support for programs in sensing system technology from all stakeholders in infrastructure security and maintenance. This can include insurance companies, engineering societies and organizations, owners, and industry, as well as national, state and local government organizations. Strategic scenarios (e.g. infrastructure security and maintenance and homeland defense) are needed in order to get practical support from the government and public for a major sensing system initiative.

Again, if these devices are developed in isolation, nobody will use them. This workshop revealed that every aspect of sensor research--from device design to application development-- requires the skills of a variety of individuals if any real world value is to be expected. For example, the development and optimal integration of a wireless sensor takes knowledge of RF, networking, MEMS design, programming, and whatever field of study that the sensor data is expected to inform. No one piece is more important than the other.

With the traditional model of funding, the pieces can be completed individually but they will exist in distant parts of the world, or different buildings on a university campus, and will not result in a system that's truly useful. Another area ripe for research is the synthesis of modeling interpretation with sensing. Think of the synergistic research benefits if the modeler collaborates with someone who actually takes measurements. The experimentalist would have a better idea what needs to be measured to be useful to develope a model, and the modeler would have a more accurate understanding of what constitutes a model that could work in the real world. Multi-disciplinary teams of dedicated researchers working toward a shared goal are the key to bringing the future sensor networks described in this report from the lab bench into the real world.

5.2. Research initiatives

Inside the US, champions within many government and industry organizations are needed in order to support a major initiative in sensing systems: DOT, DOE, DOD, GSA, and NASA. Test beds can bring

developers and users together, and demonstrate the feasibility and the economics of new sensing systems. An effective test bed should demonstrate a practical and complete sensing network, including robust wireless communications and long-lived power sources. These test beds should address challenging, meaningful, but achievable problems. On a larger scale, strategic scenarios (e.g. infrastructure security and maintenance, homeland defense) are needed in order to get practical support from the government and public for a major sensing system initiative. Many real-world sensing system problems involve retrofit of existing structures, machinery and equipment, and technologies suitable for retrofit should be emphasized.

Problems of maintaining and securing infrastructure are common between the US, Europe and Japan. International cooperation would be valuable, and should be encouraged. Government agencies that could be involved include the European Science Foundation, and METI (Ministry of Economy, Trade, and Industry in Japan). Several joint U.S. - Japan sensor meetings also began to take shape during the workshop.

5.3. Role of the private sector

Although the NSF focuses on university research, the private sector plays a vital role in sensor research and development. In fact, sensors would not be widely and inexpensively available if they were not manufactured by the private sector. Both small businesses and large industries are key to sensor proliferation. Small businesses are often closely tied to academic researchers, and are agile enough to take new ideas to market quickly. Examples include Crossbow, Dust Inc., and MicroStrain. Given the enormous capital investment required to build and operate a MEMS fabrication facility, the strength and power of large industry is often required to bring research to fruition. Examples in the MEMS-based sensor field include Honeywell, Analog Devices, and Delphi. The workshop featured a panel discussion with representatives of small businesses, large industries, and venture capitalists.

Most of the academic research in sensors involves fabrication or novel designs to solve a particular problem not solved by a commercial product. Often though, the unique sensors designed in academia could be of use to other researchers and engineers. Most small business start-ups in sensors grow out of these kinds of situations. Several participants who discussed the business of sensor networks outlined the importance of identifying potential products and possible markets. A distinction was made between a dream and a pragmatic, realistic path for commercialization.

A successful business plan has a thought-out focus and path where the end goal is helping a client solve a real, not merely perceived, problem. Often, the simplest solution is the best, most cost effective, and marketable. The speakers forcefully emphasized that a start-up should focus on specific client needs. Casting a net over too wide of a market leads to lack of focus and a hopeless dilution of resources. A successful business plan identifies the customer base as well as the technology base. It should always be remembered that only about 10% of high technology start-ups are successful. For those that do succeed, the biggest challenge becomes growth and stability over the long-term.

At the workshop, two suggested approaches to financing new firms emerged. One is to develop an idea or product and then try to acquire funding from venture capitalists or corporate incubators. The other is to start out very lean, working hard for little remuneration, and plow all proceeds back into building the company. Each approach has positives and negatives. A good VC can bring advanced business skills to the table along with important relationships throughout the industry. They can supply an impartial point of view. These services come at a cost though, often much more than half the company's stock and the loss of IP rights by the inventors. Independently bootstrapping a firm involves

tremendous sacrifices by the founders and employees, and might not be an option for innovators with, for example, family responsibilities.

The workshop participants came to broad agreement that there needs to be a smoothly operating and realistic system whereby the research migrates from the university to industry (i.e. fund research, identify start-ups, or integrate with existing applications). One way is to promote "people exchanges" between academia and industry (internships for faculty) to have more effective technology transfer, but a mechanism must be found to make the exchanges more practical and beyond NSF-GOALI and SBIR types of programs.

5.4. Success of workshop format

We believe that the workshop was a success for many reasons. Perhaps most important is the way researchers from very different backgrounds worked and communicated so well with each other. If it can be done at a workshop, it can be done at an NSF proposal review panel. As can be gleaned from the comments available on-line at http://www.ce.berkeley.edu/xxx, the participants thought the workshop was a tremendous learning experience. This single event has led to the founding of new companies and submission of proposals that did not exist before the workshop.

The workshop provided an opportunity for researchers from different countries to converge in a relaxed and informative environment. For example, there was a strong interaction between NSF program directors and their counterparts from Europe, hopefully leading to joint research programs. There was real dialogue between industrial and academic researchers, with both sides of the research coin gaining insight into the needs and realities of the other. For the organizers, perhaps the defining moment came on Wednesday during the presentation of the fifth breakout group. The presenter, who had stated on numerous occasions that he had no interest in civil and mechanical engineering, presented a thorough and thoughtful discussion of structural health monitoring and prognostication of a bridge's stability. The presentation was given with feeling and passion.

Acknowledgements

Appendix 1. List of abbreviations

AFOSR	: air force office of scientific research
CMOS	: complementary metal oxide semiconductor
CMS	: division of civil and mechanical systems, national science foundation
COTS	: commercial-off-the-shelf
DARPA	: defence advance research administration
DOE	: department of energy
DOT	: department of transportation
FRP	: fiber-reinforced polymer
GSA	: general service administration

HVAC   : heating, ventilation, and air conditioning
IC   : integrated circuit
IP   : intellectual property
IT   : information technology
JPL   : jet propulsion laboratory
MCU   : micro-controller unit
MEMS   : micro-electro-mechanical systems
METI   : ministry of economy, trade and industry, (government of Japan)
NASA   : national aeronautics and space administration
NIH   : national institutes of health
NIST   : national institute of standards and technology
NSF   : national science foundation
RF   : radio frequency
R&D   : research and development

CC