Feature

Wireless Sensor Networks: Applications and Challenges of Ubiquitous Sensing

Daniele Puccinelli and Martin Haenggi

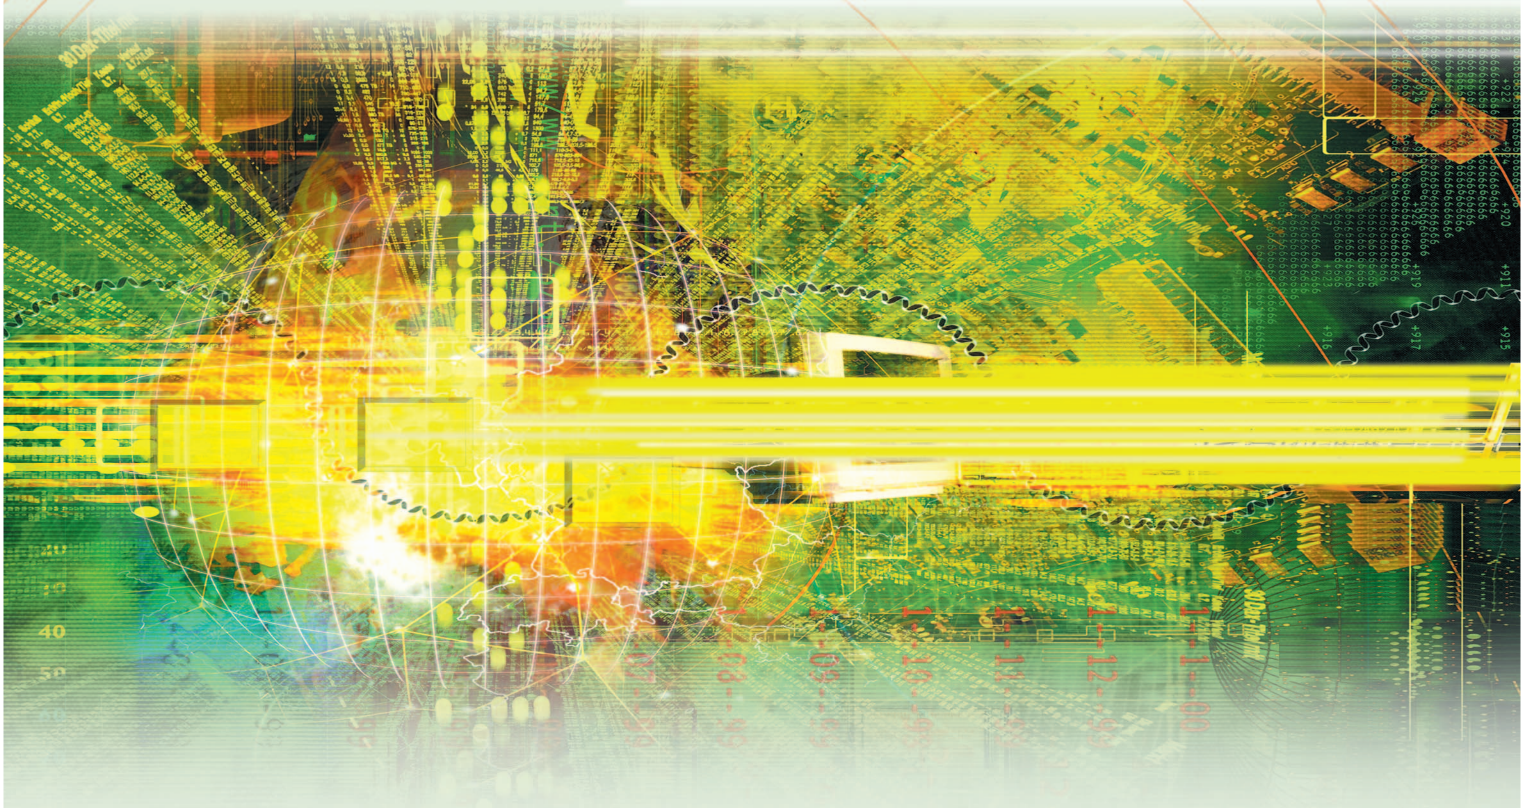

© DIGITAL VISION

Abstract

Sensor networks offer a powerful combination of distributed sensing, computing and communication. They lend themselves to countless applications and, at the same time, offer numerous challenges due to their peculiarities, primarily the stringent energy constraints to which sensing nodes are typically subjected. The distinguishing traits of sensor networks have a direct impact on the hardware design of the nodes at at least four levels: power source, processor, communication hardware, and sensors. Various hardware platforms have already been designed to test the many ideas spawned by the research community and to implement applications to virtually all fields of science and technology. We are convinced that CAS will be able to provide a substantial contribution to the development of this exciting field.

1531-636X/05/$20.00©2005 IEEE

1. Introduction

The increasing interest in wireless sensor networks can be promptly understood simply by thinking about what they essentially are: a large number of small sensing self-powered nodes which gather information or detect special events and communicate in a wireless fashion, with the end goal of handing their processed data to a base station. Sensing, processing and communication are three key elements whose combination in one tiny device gives rise to a vast number of applications [A1], [A2]. Sensor networks provide endless opportunities, but at the same time pose formidable challenges, such as the fact that energy is a scarce and usually non-renewable resource. However, recent advances in low power VLSI, embedded computing, communication hardware, and in general, the convergence of computing and communications, are making this emerging technology a reality [A3]. Likewise, advances in nanotechnology and Micro Electro-Mechanical Systems (MEMS) are pushing toward networks of tiny distributed sensors and actuators.

2. Applications of Sensor Networks

Possible applications of sensor networks are of interest to the most diverse fields. Environmental monitoring, warfare, child education, surveillance, micro-surgery, and agriculture are only a few examples [A4]. Through joint efforts of the University of California at Berkeley and the College of the Atlantic, environmental monitoring is carried out off the coast of Maine on Great Duck Island by means of a network of *Berkeley motes* equipped with various sensors [B6]. The nodes send their data to a base station which makes them available on the Internet. Since habitat monitoring is rather sensitive to human presence, the deployment of a sensor network provides a non-invasive approach and a remarkable degree of granularity in data acquisition [B7]. The same idea lies behind the Pods project at the University of Hawaii at Manoa [B8], where environmental data (air temperature, light, wind, relative humidity and rainfall) are gathered by a network of weather sensors embedded in the communication units deployed in the South-West Rift Zone in Volcanoes National Park on the Big Island of Hawaii. A major concern of the researchers was in this case camouflaging the sensors to make them invisible to curious tourists. In Princeton's Zebranet Project [B9], a dynamic sensor network has been created by attaching special collars equipped with a low-power GPS system to the necks of zebras to monitor their moves and their behavior. Since the network is designed to operate in an infrastructure-free environment, peer-to-peer swaps of information are used to produce redundant databases so that researchers only have to encounter a few zebras in order to collect the data. Sensor networks can also be used to monitor and study natural phenomena which intrinsically discourage human presence, such as hurricanes and forest fires. Joint efforts between Harvard University, the University of New Hampshire, and the University of North Carolina have recently led to the deployment of a wireless sensor network to monitor eruptions at Volcán Tungurahua, an active volcano in central Ecuador. A network of Berkeley motes monitored infrasonic signals during eruptions, and data were transmitted over a 9 km wireless link to a base station at the volcano observatory [B10].

Intel's Wireless Vineyard [B11] is an example of using ubiquitous computing for agricultural monitoring. In this application, the network is expected not only to collect and interpret data, but also to use such data to make decisions aimed at detecting the presence of parasites and enabling the use of the appropriate kind of insecticide. Data collection relies on *data mules*, small devices carried by people (or dogs) that communicate with the nodes and collect data. In this project, the attention is shifted from reliable information collection to active decision-making based on acquired data.

Just as they can be used to monitor nature, sensor networks can likewise be used to monitor human behavior. In the Smart Kindergarten project at UCLA [B12], wirelessly-networked, sensor-enhanced toys and other classroom objects supervise the learning process of children and allow unobtrusive monitoring by the teacher.

Medical research and healthcare can greatly benefit from sensor networks: vital sign monitoring and accident recognition are the most natural applications. An important issue is the care of the elderly, especially if they are affected by cognitive decline: a network of sensors and actuators could monitor them and even assist them in their daily routine. Smart appliances could help them organize their lives by reminding them of their meals and medications. Sensors can be used to capture vital signs from patients in real-time and relay the data to handheld computers carried by medical personnel, and wearable sensor nodes can store patient data such as identification, history, and treatments. With these ideas in mind, Harvard University is cooperating with the School of Medicine at Boston University to develop CodeBlue, an infrastructure designed to support wireless medical sensors, PDAs, PCs, and other devices that may be used to monitor and treat patients in various medical scenarios [B13]. On the hardware side, the research team has

---

Martin Haenggi is with the Department of Electrical Engineering, University of Notre Dame, Notre Dame, IN 46556; Fax +1 574 631 4393; mhaenggi@@nd.edu. Daniele Puccinelli is also with the Department of Electrical Engineering, University of Notre Dame, Notre Dame, IN 46556.

created Vital Dust, a set of devices based on the MICA2[1] sensor node platform (one of the most popular members of the Berkeley motes family), which collect heart rate, oxygen saturation, and EKG data and relay them over a medium-range (100 m) wireless network to a PDA [B14]. Interactions between sensor networks and humans are already judged controversial. The US has recently approved the use of a radio-frequency implantable device (VeriChip) on humans, whose intended application is accessing the medical records of a patient in an emergency. Potential future repercussions of this decision have been discussed in the media.

An interesting application to civil engineering is the idea of Smart Buildings: wireless sensor and actuator networks integrated within buildings could allow distributed monitoring and control, improving living conditions and reducing the energy consumption, for instance by controlling temperature and air flow.

Military applications are plentiful. An intriguing example is DARPA's self-healing minefield [B15], a self-organizing sensor network where peer-to-peer communication between anti-tank mines is used to respond to attacks and redistribute the mines in order to heal breaches, complicating the progress of enemy troops. Urban warfare is another application that distributed sensing lends itself to. An ensemble of nodes could be deployed in a urban landscape to detect chemical attacks, or track enemy movements. PinPtr is an ad hoc acoustic sensor network for sniper localization developed at Vanderbilt University [B16]. The network detects the muzzle blast and the acoustic shock wave that originate from the sound of gunfire. The arrival times of the acoustic events at different sensor nodes are used to estimate the position of the sniper and send it to the base station with a special data aggregation and routing service.

Going back to peaceful applications, efforts are underway at Carnegie Mellon University and Intel for the design of IrisNet (Internet-scale Resource-Intensive Sensor Network Services) [B17], an architecture for a worldwide sensor web based on common computing hardware such as Internet-connected PCs and low-cost sensing hardware such as webcams. The network interface of a PC indeed senses the virtual environment of a LAN or the Internet rather than a physical environment; with an architecture based on the concept of a distributed database [B18], this hardware can be orchestrated into a global sensor system that responds to queries from users.

3. Characteristic Features of Sensor Networks

In ad hoc networks, wireless nodes self-organize into an infrastructureless network with a dynamic topology. Sensor networks (such as the one in Figure 1) share these traits, but also have several distinguishing features. The number of nodes in a typical sensor network is much higher than in a typical ad hoc network, and dense deployments are often desired to ensure coverage and connectivity; for these reasons, sensor network hardware must be cheap. Nodes typically have stringent energy limitations, which make them more failure-prone. They are generally assumed to be stationary, but their relatively frequent breakdowns and the volatile nature of the wireless channel nonetheless result in a variable network topology. Ideally, sensor network hardware should be power-efficient, small, inexpensive, and reliable in order to maximize network lifetime, add flexibility, facilitate data collection and minimize the need for maintenance.

Lifetime

Lifetime is extremely critical for most applications, and its primary limiting factor is the energy consumption of the nodes, which need to be self-powering. Although it is often assumed that the transmit power associated with packet transmission accounts for the lion's share of power consumption, sensing, signal processing and even hardware operation in standby mode consume a consistent amount of power as well [C19], [C20]. In some applications, extra power is needed for macro-scale actuation.

Many researchers suggest that energy consumption could be reduced by considering the existing interdependencies between individual layers in the network protocol stack. Routing and channel access protocols, for instance, could greatly benefit from an information exchange with the physical layer.

At the physical layer, benefits can be obtained with lower radio duty cycles and dynamic modulation scaling (varying the constellation size to minimize energy expenditure

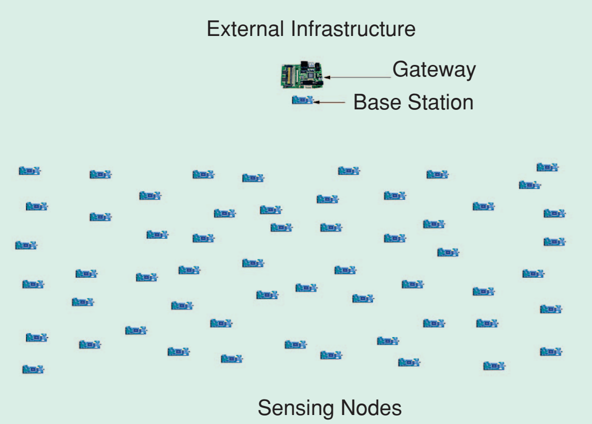

Figure 1. A generic sensor network with a two-tiered architecture.

---

[1] See Section 5 for a hardware overview.

[D35]). Using low-power modi for the processor or disabling the radio is generally advantageous, even though periodically turning a subsystem on and off may be more costly than always keeping it on. Techniques aimed at reducing the idle mode leakage current in CMOS-based processors are also noteworthy [D36].

Medium Access Control (MAC) solutions have a direct impact on energy consumption, as some of the primary causes of energy waste are found at the MAC layer: collisions, control packet overhead and idle listening. Power-saving forward error control techniques are not easy to implement due to the high amount of computing power that they require and the fact that long packets are normally not practical.

Energy-efficient routing should avoid the loss of a node due to battery depletion. Many proposed protocols tend to minimize energy consumption on forwarding paths, but if some nodes happen to be located on most forwarding paths (*e.g.*, close to the base station), their lifetime will be reduced.

Flexibility

Sensor networks should be scalable, and they should be able to dynamically adapt to changes in node density and topology, like in the case of the self-healing minefields. In surveillance applications, most nodes may remain quiescent as long as nothing interesting happens. However, they must be able to respond to special events that the network intends to study with some degree of granularity. In a self-healing minefield, a number of sensing mines may sleep as long as none of their peers explodes, but need to quickly become operational in the case of an enemy attack. Response time is also very critical in control applications (sensor/actuator networks) in which the network is to provide a delay-guaranteed service.

Untethered systems need to self-configure and adapt to different conditions. Sensor networks should also be robust to changes in their topology, for instance due to the failure of individual nodes. In particular, connectivity and coverage should always be guaranteed. Connectivity is achieved if the base station can be reached from any node. Coverage can be seen as a measure of quality of service in a sensor network [C23], as it defines how well a particular area can be observed by a network and characterizes the probability of detection of geographically constrained phenomena or events. Complete coverage is particularly important for surveillance applications.

Maintenance

The only desired form of maintenance in a sensor network is the complete or partial update of the program code in the sensor nodes over the wireless channel. All sensor nodes should be updated, and the restrictions on the size of the new code should be the same as in the case of wired programming. Packet loss must be accounted for and should not impede correct reprogramming. The portion of code always running in the node to guarantee reprogramming support should have a small footprint, and updating procedures should only cause a brief interruption of the normal operation of the node [C24].

The functioning of the network as a whole should not be endangered by unavoidable failures of single nodes, which may occur for a number of reasons, from battery depletion to unpredictable external events, and may either be independent or spatially correlated [C25]. Fault-tolerance is particularly crucial as ongoing maintenance is rarely an option in sensor network applications.

Self-configuring nodes are necessary to allow the deployment process to run smoothly without human interaction, which should in principle be limited to placing nodes into a given geographical area. It is not desirable to have humans configure nodes for habitat monitoring and destructively interfere with wildlife in the process, or configure nodes for urban warfare monitoring in a hostile environment. The nodes should be able to assess the quality of the network deployment and indicate any problems that may arise, as well as adjust to changing environmental conditions by automatic reconfiguration. Location awareness is important for self-configuration and has definite advantages in terms of routing [C26] and security. Time synchronization [C27] is advantageous in promoting cooperation among nodes, such as data fusion, channel access, coordination of sleep modi, or security-related interaction.

Data Collection

Data collection is related to network connectivity and coverage. An interesting solution is the use of ubiquitous mobile agents that randomly move around to gather data bridging sensor nodes and access points, whimsically named data MULEs (Mobile Ubiquitous LAN Extensions) in [C28]. The predictable mobility of the data sink can be used to save power [C29], as nodes can learn its schedule. A similar concept has been implemented in Intel's Wireless Vineyard.

It is often the case that all data are relayed to a base station, but this form of centralized data collection may shorten network lifetime. Relaying data to a data sink causes non-uniform power consumption patterns that may overburden forwarding nodes [C21]. This is particularly harsh on nodes providing end links to base stations, which may end up relaying traffic coming from all other nodes, thus forming a critical bottleneck for network throughput [A4], [C22], as shown in Figure 2.

An interesting technique is clustering [C30]: nodes team up to form clusters and transmit their information to their cluster heads, which fuse the data and forward it to a sink. Fewer packets are transmitted, and a uniform energy consumption pattern may be achieved by periodic re-clustering. Data redundancy is minimized, as the aggregation process fuses strongly correlated measurements.

Many applications require that queries be sent to sensing nodes. This is true, for example, whenever the goal is gathering data regarding a particular area where various sensors have been deployed. This is the rationale behind looking at a sensor network as a database [C31].

A sensor network should be able to protect itself and its data from external attacks, but the severe limitations of lower-end sensor node hardware make security a true challenge. Typical encryption schemes, for instance, require large amounts of memory that are unavailable in sensor nodes. Data confidentiality should be preserved by encrypting data with a secret key shared with the intended receiver. Data integrity should be ensured to prevent unauthorized data alteration. An authenticated broadcast must allow the verification of the legitimacy of data and their sender. In a number of commercial applications, a serious disservice to the user of a sensor network is compromising data availability (denial of service), which can be achieved by sleep-deprivation torture [C33]: batteries may be drained by continuous service requests or demands for legitimate but intensive tasks [C34], preventing the node from entering sleep modi.

4. Hardware Design Issues

In a generic sensor node (Figure 3), we can identify a power module, a communication block, a processing unit with internal and/or external memory, and a module for sensing and actuation.

Power

Using stored energy or harvesting energy from the outside world are the two options for the power module. Energy storage may be achieved with the use of batteries or alternative devices such as fuel cells or miniaturized heat engines, whereas energy-scavenging opportunities [D37] are provided by solar power, vibrations, acoustic noise, and piezoelectric effects [D38]. The vast majority of the existing commercial and research platforms relies on batteries, which dominate the node size. Primary (non-rechargeable) batteries are often chosen, predominantly AA, AAA and coin-type. Alkaline batteries offer a high energy density at a cheap price, offset by a non-flat discharge, a large physical size with respect to a typical sensor node, and a shelf life of only 5 years. Voltage regulation could in principle be employed, but its high inefficiency and large quiescent current consumption call for the use of components that can deal with large variations in the supply voltage [A5]. Lithium cells are very compact and boast a flat discharge curve. Secondary (rechargeable) batteries are typically not desirable, as they offer a lower energy density and a higher cost, not to mention the fact that in most applications recharging is simply not practical.

Fuel cells [D39] are rechargeable electrochemical energy-conversion devices where electricity and heat are produced as long as hydrogen is supplied to react with oxygen. Pollution is minimal, as water is the main byproduct of the reaction. The potential of fuel cells for energy storage and power delivery is much higher than the one of traditional battery technologies, but the fact that they require hydrogen complicates their application. Using renewable energy and scavenging techniques is an interesting alternative.

Communication

Most sensor networks use radio communication, even if alternative solutions are offered by laser and infrared. Nearly all radio-based platforms use COTS (Commercial Off-The-Shelf) components. Popular choices include the TR1000 from RFM (used in the MICA motes) and the CC1000 from Chipcon (chosen for the MICA2 platform). More recent solutions use industry standards like IEEE 802.15.4 (MICAz and Telos motes with CC2420 from Chipcon) or pseudo-standards like Bluetooth. Typically, the transmit power ranges between −25 dBm and 10 dBm, while the receiver sensitivity can be as good as −110 dBm.

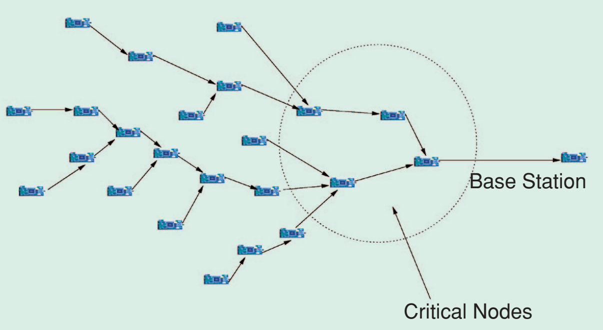

Figure 2. A uniform energy consumption pattern should avoid the depletion of the resources of nodes located in the vicinities of the base station.

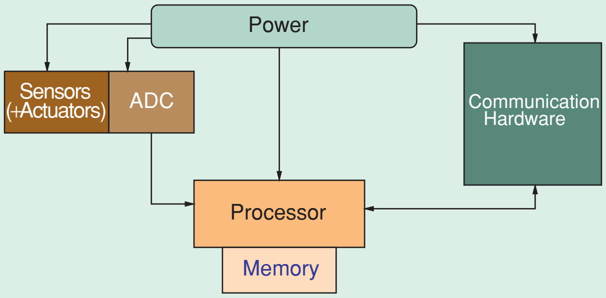

Figure 3. Anatomy of a generic sensor node.

Spread spectrum techniques increase the channel reliability and the noise tolerance by spreading the signal over a wide range of frequencies. Frequency hopping (FH) is a spread spectrum technique used by Bluetooth: the carrier frequency changes 1600 times per second on the basis of a pseudo-random algorithm. However, channel synchronization, hopping sequence search, and the high data rate increase power consumption; this is one of the strongest caveats when using Bluetooth in sensor network nodes. In Direct Sequence Spread Spectrum (DSSS), communication is carried out on a single carrier frequency. The signal is multiplied by a higher rate pseudo-random sequence and thus spread over a wide frequency range (typical DSSS radios have spreading factors between 15 and 100).

Ultra Wide Band (UWB) is of great interest for sensor networks since it meets some of their main requirements. UWB is a particular carrier-free spread spectrum technique where the RF signal is spread over a spectrum as large as several GHz. This implies that UWB signals look like noise to conventional radios. Such signals are produced using baseband pulses (for instance, Gaussian monopulses) whose length ranges from 100 ps to 1 ns, and baseband transmission is generally carried out by means of pulse position modulation (PPM). Modulation and demodulation are indeed extremely cheap. UWB provides built-in ranging capabilities (a wideband signal allows a good time resolution and therefore a good location accuracy) [D40], allows a very low power consumption, and performs well in the presence of multipath fading.

Radios with relatively low bit-rates (up to 100 kbps) are advantageous in terms of power consumption. In most sensor networks, high data rates are not needed, even though they allow shorter transmission times thus permitting lower duty cycles and alleviating channel access contention. It is also desirable for a radio to quickly switch from a sleep mode to an operational mode.

Optical transceivers such as lasers offer a strong power advantage, mainly due to their high directionality and the fact that only baseband processing is required. Also, security is intrinsically guaranteed (intercepted signals are altered). However, the need for a line of sight and precise localization makes this option impractical for most applications.

Processing and Computing

Although low-power FPGAs might become a viable option in the near future [D41], microcontrollers (MCUs) are now the primary choice for processing in sensor nodes. The key metric in the selection of an MCU is power consumption. Sleep modi deserve special attention, as in many applications low duty cycles are essential for lifetime extension. Just as in the case of the radio module, a fast wake-up time is important. Most CPUs used in lower-end sensor nodes have clock speeds of a few MHz. The memory requirements depend on the application and the network topology: data storage is not critical if data are often relayed to a base station. Berkeley motes, UCLA's Medusa MK-2 and ETHZ's BTnodes use low-cost Atmel AVR 8-bit RISC microcontrollers which consume about 1500 pJ/instruction. More sophisticated platforms, such as the Intel iMote and Rockwell WINS nodes, use Intel StrongArm/XScale 32-bit processors.

Sensing

The high sampling rates of modern digital sensors are usually not needed in sensor networks. The power efficiency of sensors and their turn-on and turn-off time are much more important. Additional issues are the physical size of the sensing hardware, fabrication, and assembly compatibility with other components of the system. Packaging requirements come into play, for instance, with chemical sensors which require contact with the environment [D42]. Using a microcontroller with an on-chip analog comparator is another energy-saving technique which allows the node to avoid sampling values falling outside a certain range [D43]. The ADC which complements analog sensors is particularly critical, as its resolution has a direct impact on energy consumption. Fortunately, typical sensor network applications do not have stringent resolution requirements.

Micromachining techniques have allowed the miniaturization of many types of sensors. Performance does decrease with sensor size, but for many sensor network applications size matters much more than accuracy.

Standard integrated circuits may also be used as temperature sensors (*e.g.*, using the temperature-dependence of subthreshold MOSFETs and pn junctions) or light intensity transducers (*e.g.*, using photodiodes or phototransistors) [D44]. Nanosensors can offer promising solutions for biological and chemical sensors while concurrently meeting the most ambitious miniaturization needs.

5. Existing Hardware Platforms

Berkeley motes, made commercially available by Crossbow, are by all means the best known sensor node hardware implementation, used by more than 100 research organizations. They consist of an embedded microcontroller, low-power radio, and a small memory, and they are powered by two AA batteries. MICA and MICA2 are the most successful families of Berkeley motes. The MICA2 platform, whose layout is shown in Figure 4, is equipped with an Atmel ATmega128L and has a CC1000 transceiver. A 51-pin expansion connector is available to interface sensors (commercial sensor boards designed for this specific platform are available). Since the MCU is to handle medium access and baseband processing, a fine-grained event-driven real-time operating system (TinyOS) has been implemented to specifically address the concurrency and resource management needs of sensor nodes. For applications that require a better form factor, the circular MICA2Dot can be used: it has most of the resources of MICA2, but is only 2.5 cm in diameter. Berkeley motes up to the MICA2 generation cannot interface with other wireless-enabled devices [E47]. However, the newer generations MICAz and Telos support IEEE 802.15.4, which is part of the 802.15 Wireless Personal Area Network (WPAN) standard being developed by IEEE. At this point, these devices represent a very good solution for generic sensing nodes, even though their unit cost is still relatively high (about $100–$200). The proliferation of different lower-end hardware platforms within the Berkeley mote family has recently led to the development of a new version of TinyOS which introduces a flexible hardware abstraction architecture to simplify multi-platform support [E48].

Tables 1 and 2 show an overview of the radio transceivers and the microcontrollers most commonly used in existing hardware platforms; an overview of the key features of the platforms is provided in Table 3.

Intel has designed its own iMote [E49] to implement various improvements over available mote designs, such as increased CPU processing power, increased main memory size for on-board computing and improved radio reliability. In the iMote, a powerful ARM7TDMI core is complemented by a large main memory and non-volatile storage area; on the radio side, Bluetooth has been chosen.

Various platforms have been developed for the use of Berkeley motes in mobile sensor networks to enable investigations into controlled mobility, which facilitates deployment and network repair and provides possibilities for the implementation of energy-harvesting. UCLA's RoboMote [E50], Notre Dame's MicaBot [E51] and UC Berkeley's CotsBots [E52] are examples of efforts in this direction.

UCLA's Medusa MK-2 sensor nodes [E53], developed for the Smart Kindergarten project, expand Berkeley motes with a second microcontroller. An on-board power management and tracking unit monitors power consumption within the different subsystems and selectively powers down unused parts of the node.

UCLA has also developed iBadge [E54], a wearable sensor node with sufficient computational power to process the sensed data. Built around an ATMega128L and a DSP, it features a Localization Unit designed to estimate the position of iBadge in a room based on the presence of special nodes of known location attached to the ceilings.

In the context of the EYES project (a joint effort among several European institutions) custom nodes [E55], [C24] have been developed to test and demonstrate energy-efficient networking algorithms. On the software side, a proprietary operating system, PEEROS (Preemptive EYES Real Time Operating System), has been implemented.

The Smart-Its project has investigated the possibility of embedding computational power into objects, leading to the creation of three hardware platforms: DIY Smart-its, Particle Computers and BTnodes.

The DIY Smart-its [E56] have been developed in the UK at Lancaster University; their modular design is based on a core board that provides processing and communication and can be extended with add-on boards. A typical setup of Smart-its consists of one or more sensing nodes that broadcast their data to a base station which consists of a standard core board connected to the serial port of a PC. Simplicity and extensibility are the key features of this platform, which has been developed for the creation of Smart Objects. An interesting application is the Weight Table: four load cells placed underneath a coffee table form a Wheatstone bridge and are connected to a DIY node that observes load changes, determines event types like placement and removal of objects or a person moving a finger across the surface, and also retrieves the position of an object by correlating the values of the individual load cells after the event type (removed or placed) has been recognized [E57].

Particle Computers have been developed at the University of Karlsruhe, Germany. Similarly to the DIY platform, the Particle Smart-its are based on a core board equipped with a Microchip PIC; they are optimized for energy efficiency, scalable communication and small scale (17 mm × 30 mm). Particles communicate in an ad hoc fashion: as two Particles come close to one another,

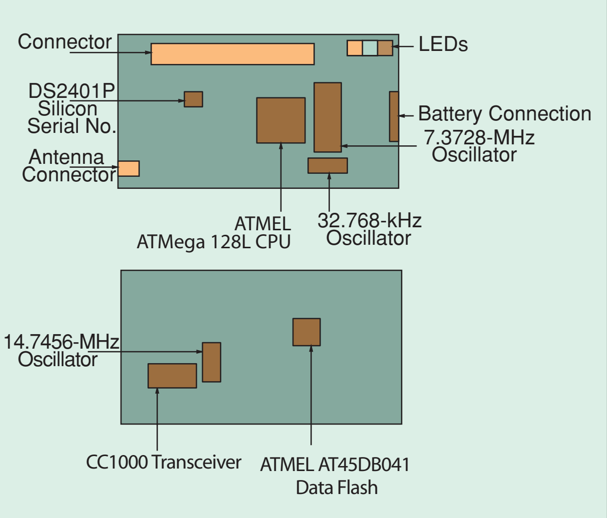

Figure 4. Layout of the MICA2 platform.

they are able to talk. Additionally, if Particles come near a gateway device, they can be connected to Internet-enabled devices and access services and information on the Internet as well as provide information [E58].

The BTnode hardware from ETHZ [E47] is based on an Atmel ATmega128L microcontroller and a Bluetooth module. Although advertised as a low-power technology, Bluetooth has a relatively high power consumption, as discussed before. It also has long connection setup times and a lower degree of freedom with respect to possible network topologies. On the other hand, it ensures interoperability between different devices, enables application development through a standardized interface, and offers a significantly higher bandwidth (about 1 Mbps) compared to many low-power radios (about 50 Kbps). Moreover, Bluetooth support means that COTS hardware can be used to create a gateway between a sensor network and an external network (*e.g.*, the Internet), as opposed to more costly proprietary solutions [E59].

MIT is working on the μAMPS (μ-Adaptive Multi-domain Power-aware Sensors) project, which explores energy-efficiency constraints and key issues such as self-configuration, reconfigurability, and flexibility. A first prototype has been designed with COTS components: three stackable boards (processing, radio and power) and an optional extension module. The energy dissipation of this microsensor node is reduced through a variety of power-aware design techniques [D45] including fine-grain shutdown of inactive components, dynamic voltage and frequency scaling of the processor core, and adjustable radio transmission power based on the required range. Dynamic voltage scaling is a technique used for active power management where the supply voltage and clock frequency of the processor are regulated depending on the computational load, which can vary significantly based on the operational mode [D36], [C20]. The main goal of second generation μAMPS is clearly stated in [D46] as breaking the 100 μW average power barrier.

Another interesting MIT project is the Pushpin computing system [E60], whose goal is the modelling, testing, and deployment of distributed peer-to-peer sensor networks consisting of many identical nodes. The pushpins are 18 mm × 18 mm modular devices with a power substrate, an infrared communication module, a processing module (Cygnal C8051F016) and an expansion module (*e.g.*, for sensors); they are powered by direct contact between the power substrate and layered conductive sheets.

Table 1.
Radios used in sensor node platforms.

Radio (Manufacturer)	Band [MHz]	Max. Data Rate [kbps]	Sensit. [dBm]	Notes
TR1000 (RFM)	916.5	115.2	−106	OOK/ASK
TR1001 (RFM)	868.35	115.2	−106	OOK/ASK
CC1000 (Chipcon)	300–1,000	76.8	−110	FSK, −20 to 10 dBm
CC2420 (Chipcon)	2,400	250	−94	OQPSK, −24 to 0 dBm, IEEE 802.15.4, DSSS
BiM2 (Radiometrix)	433.92	64	−93	
9XStream (MaxStream)	902–928	20	−114	FHSS

Table 2.
Microcontrollers used in sensor node platforms.

MCU	Max. Freq. [MHz]	Memory	Data Size [bits]	ADC [bits]	Architecture
AT90LS8535 (Atmel)	4	8 kB Flash, 512B EEPROM, 512B SRAM	8	10	AVR
ATMega128L (Atmel)	8	128 kB Flash, 4 kB EEPROM, 4 kB SRAM	8	10	AVR
AT91FR4081 (Atmel)	33	136 kB On-Chip SRAM, 8 Mb Flash	32	–	Based on ARM core (ARM7TDMI)
MSP430F149 (TI)	8	60 kB + 256B Flash, 2 kB RAM	16	12	Von Neumann
C8051F016 (Cygnal)	25	2304B RAM, 32 kB Flash	8	10	Harvard 8051
PIC18F6720 (Microchip)	25	128 kB Flash, 3840B SRAM, 1 kB EEPROM	8	10	Harvard
PIC18F252 (Microchip)	40	32 K Flash, 1536B RAM, 256B EEPROM	8	10	Harvard
StrongARM SA-1110 (Intel)	133	–	32	–	ARM v.4
PXA255 (Intel)	400	32 kB Instruction Cache, 32 kB Data Cache, 2 kB *Mini* Data Cache	32	–	ARM v.5TE

MIT has also built Tribble (Tactile reactive interface built by linked elements), a spherical robot wrapped by a wired skinlike sensor network designed to emulate the functionalities of biological skin [E61]. Tribble's surface is divided into 32 patches with a Pushpin processing module and an array of sensors and actuators.

At Lancaster University, surfaces provide power and network connectivity in the Pin&Play project. Network nodes come in different form factors, but all share the Pin&Play connector, a custom component that allows physical connection and networking through conductive sheets which are embedded in surfaces such as a wall or a bulletin board [E62]. Pin&Play falls in between wired and wireless technologies as it provides network access and power across 2D surfaces. Wall-mounted objects are especially suited to be augmented to become Pin&Play objects. In a demonstration, a wall switch was augmented and freely placed anywhere on a wall with a Pin&Play surface as wallpaper.

For applications which do not call for the minimization of power consumption, high-end nodes are available. Rockwellís WINS nodes and Sensoria's WINS 3.0 Wireless Sensing Platform are equipped with more powerful processors and radio systems. The embedded PC modules based on widely supported standards PC/104 and PC/104-plus feature Pentium processors; moreover, PC/104 peripherals include digital I/O devices, sensors and actuators, and PC-104 products support almost all PC software. PFU Systems' Plug-N-Run products, which feature Pentium processors, also belong to this category. They offer the capabilities of PCs and the size of a sensor node, but lack built-in communication hardware. COTS components or lower-end nodes may be used in this sense [C32]. Research is underway toward the creation of sensor nodes that are more capable than the motes, yet smaller and more power-efficient than higher-end nodes.

Simple yet effective gateway devices are the MIB programming boards from Crossbow, which bridge networks of Berkeley motes with a PC (to which they interface using the serial port or Ethernet). In the case of Telos motes, any generic node (*i.e.*, any Telos mote) can act as a gateway, as it may be connected to the USB port of a PC and bridge it to the network. Of course, more powerful gateway devices are also available. Crossbow's Stargate is a powerful embedded computing platform (running Linux) with enhanced communication and sensor signal processing capabilities based on Intel PXA255, the same X-Scale processor that forms the core of Sensoria WINS 3.0 nodes. Stargate has a connector for Berkeley motes, may be bridged to a PC via Ethernet or 802.11, and includes built-in Bluetooth support.

6. Closing Remarks

Sensor networks offer countless challenges, but their versatility and their broad range of applications are eliciting more and more interest from the research community as well as from industry. Sensor networks have the potential of triggering the next revolution in information technology. The challenges in terms of circuits and systems are numerous: the development of low-power communication hardware, low-power microcontrollers, MEMS-based sensors and actuators, efficient AD conversion, and energy-scavenging devices is necessary to enhance the potential and the performance of sensor networks. System integration is another major challenge that sensor networks offer to the circuits and systems research community. We believe that CAS can and should have a significant impact in this emerging, exciting area.

Table 3.
Hardware features of various platforms.

Platform	CPU	Comm.	External Memory	Power Supply
WesC (UCB)	AT90LS8535	TR1000	32 kB Flash	Lithium Battery
MICA (UCB, Xbow)	ATMega128L	TR1000	512 kB Flash	AA
MICA2 (UCB, Xbow)	ATMega128L	CC1000	512 kB Flash	AA
MICA2Dot (UCB, Xbow)	ATMega128L	CC1000	512 kB Flash	Lithium Battery
MICAz (UCB, Xbow)	ATMega128L	CC2420	512 kB Flash	AA
Telos (Moteiv)	MSP430F149	CC2420	512 kB Flash	AA
iMote (Intel)	ARM7TDMI Core	Bluetooth	64 kB SRAM, 512 kB Flash	AA
Medusa MK-2 (UCLA)	ATMega103L AT91FR4081	TR1000	4 Mb Flash	Rechargeable Lithium Ion
iBadge (UCLA)	ATMega128L	Bluetooth, TR1000	4 Mb Flash	Rechargeable Lithium Ion
DIY (Lancaster University)	PIC18F252	BiM2	64 Kb FRAM	AAA, Lithium, Rechargeable
Particle (TH)	PIC18F6720	RFM TR1001	32 kB EEPROM	AAA or Lithium Coin Battery or Rechargeable
BT Nodes (ETHZ)	ATMega128L	Bluetooth, CC1000	244 kB SRAM	AA
ZebraNet (Princeton)	MSP430F149	9XStream	4 Mb Flash	Lithium Ion
Pushpin (MIT)	C8051F016	Infrared	–	Power Substrate
WINS 3.0 (Sensoria)	PXA255	802.11b	64 MB SDRAM, 32 MB + 1 GB Flash	Batteries

Acknowledgments

The support of NSF (grants ECS 03-29766 and CAREER CNS 04-47869) is gratefully acknowledged.

References

General References

[A1] I.F. Akyildiz, W. Su, Y. Sankarasubramaniam, and E. Cayirci, "A survey on sensor networks," in *IEEE Communications Magazine*, pp. 102–114, Aug. 2002.

[A2] L.B. Ruiz, L.H.A. Correia, L.F.M. Vieira, D.F. Macedo, E.F. Nakamura, C.M.S. Figueiredo, M.A.M. Vieira, E.H.B. Maia, D. Câmara, A.A.F. Loureiro, J.M.S. Nogueira, D.C. da Silva Jr., and A.O. Fernandes, "Architectures for wireless sensor networks (In Portuguese)," in *Proceedings of the 22nd Brazilian Symposium on Computer Networks (SBRC'04)*, Gramado, Brazil, pp. 167–218, May 2004. Tutorial. ISBN: 85-88442-82-5.

[A3] C.Y. Chong and S.P. Kumar, "Sensor networks: Evolution, opportunities, and challenges," in *IEEE Proceedings*, pp. 1247–1254, Aug. 2003.

[A4] M. Haenggi, "Opportunities and Challenges in Wireless Sensor Networks," in *Handbook of Sensor Networks: Compact Wireless and Wired Sensing Systems*, M. Ilyas and I. Mahgoub, eds., Boca Raton, FL, pp. 1.1–1.14, CRC Press, 2004.

[A5] J. Hill, System Architecture for Wireless Sensor Networks. Ph.D. thesis, University of California at Berkeley, Spring 2003.

Applications

[B6] A. Mainwaring, J. Polastre, R. Szewczyk, D. Culler, and J. Anderson, "Wireless sensor networks for habitat monitoring," in *First ACM Workshop on Wireless Sensor Networks and Applications*, Atlanta, GA, Sept. 2002.

[B7] A. Cerpa, J. Elson, D. Estrin, L. Girod, M. Hamilton, and J. Zhao, "Habitat monitoring: Application driver for wireless communications technology," in *ACM SIGCOMM Workshop on Data Communications in Latin America and the Caribbean*, San José, Costa Rica, Apr. 2001.

[B8] E. Biagioni and K. Bridges, "The application of remote sensor technology to assist the recovery of rare and endangered species," *International Journal of High Performance Computing Applications*, vol. 16, pp. 315–324, Aug. 2002.

[B9] P. Juang, H. Oki, Y. Wang, M. Martonosi, L. Peh, and D. Rubenstein, "Energy-efficient computing for wildlife tracking: Design tradeoffs and early experiences with ZebraNet," in *Proceedings of the 10th International Conference on Architectural Support for Programming Languages and Operating Systems (ASPLOS-X)*, San Jose, CA, Oct. 2002.

[B10] G. Werner-Allen, J. Johnson, M. Ruiz, J. Lees, and M. Welsh, "Monitoring volcanic eruptions with a wireless sensor network," in *Proceedings of the Second European Workshop on Wireless Sensor Networks (EWSN'05)*, Jan. 2005.

[B11] J. Burrell, T. Brooke, and R. Beckwith, "Vineyard computing: Sensor networks in agricultural production," *IEEE Pervasive Computing*, vol. 3, no. 1, pp. 38–45, 2004.

[B12] M. Srivastava, R. Muntz, and M. Potkonjak, "Smart kindergarten: Sensor-based wireless networks for smart developmental problem-solving enviroments," in *Proceedings of the 7th Annual International Conference on Mobile Computing and Networking (MobiCom'01)*, Rome, Italy, pp. 132–138, 2001.

[B13] T. Fulford-Jones, D. Malan, M. Welsh, and S. Moulton, "CodeBlue: An ad hoc sensor network infrastructure for emergency medical care," in *International Workshop on Wearable and Implantable Body Sensor Networks*, London, UK, 2004.

[B14] D. Myung, B. Duncan, D. Malan, M. Welsh, M. Gaynor, and S. Moulton, "Vital dust: Wireless sensors and a sensor network for real-time patient monitoring," in *8th Annual New England Regional Trauma Conference*, Burlington, MA, 2002.

[B15] "Self-healing Mines" http://www.darpa.mil/ato/programs/SHM/.

[B16] M. Maroti, G. Simon, A. Ledeczi, and J. Sztipanovits, "Shooter localization in urban terrain," *IEEE Computer*, vol. 37, pp. 60–61, Aug. 2004.

[B17] P. Gibbons, B. Karp, Y. Ke, S. Nath, and S. Seshan, "IrisNet: An architecture for a worldwide sensor web," *IEEE Pervasive Computing*, vol. 2, no. 4, pp. 22–33, 2003.

[B18] P. Gibbons, B. Karp, Y. Ke, S. Nath, and S. Seshan, "IrisNet: An Architecture for Enabling Sensor-Enriched Internet Service," Tech. Rep. IRP-TR-03-04, Intel Research, Pittsburgh, PA, June 2003.

Characteristic Features

Lifetime

[C19] A. Goldsmith and S. Wicker, "Design challenges for energy-constrained ad hoc wireless networks," *IEEE Wireless Communications Magazine*, vol. 9, pp. 8–27, Aug. 2002.

[C20] L. Yuan and G. Qu, "Energy-efficient Design of Distributed Sensor Networks," in *Handbook of Sensor Networks: Compact Wireless and Wired Sensing Systems*, M. Ilyas and I. Mahgoub, eds., Boca Raton, FL, pp. 38.1–38.19, CRC Press, 2004.

[C21] M. Haenggi, "Twelve Reasons not to Route over Many Short Hops," in *IEEE Vehicular Technology Conference (VTC'04 Fall)*, Los Angeles, CA, Sept. 2004.

[C22] M. Haenggi, "Energy-Balancing Strategies for Wireless Sensor Networks," in *IEEE International Symposium on Circuits and Systems (ISCAS'03)*, Bangkok, Thailand, May 2003.

Coverage

[C23] S. Meguerdichian, F. Koushanfar, M. Potkonjak, and M. Srivastava, "Coverage problems in wireless ad-hoc sensor networks," in *Proceedings of the 20th Annual Joint Conference of the IEEE Computer and Communications Societies (INFOCOM'01)*, vol. 3, Anchorage, AK, pp. 1380–1387, Apr. 2001.

Maintenance

[C24] N. Reijers and K. Loangendoen, "Efficient code distribution in wireless sensor networks," in *Second ACM International Workshop on Wireless Sensor Networks and Applications*, San Diego, CA, Sept. 2003.

[C25] D. Ganesan, R. Govindan, S. Shenker, and D. Estrin, "Highly resilient, energy efficient multipath routing in wireless sensor networks," in *Proceedings of the 2nd ACM International Symposium on Mobile Ad Hoc Networking and Computing (MobiHoc'01)*, Long Beach, CA, pp. 251–254, 2001.

Localization and Synchronization

[C26] M. Mauve, H. Hartenstein, H. Fuessler, J. Widmer, and W. Effelsberg, "Positionsbasiertes Routing fuer die Kommunikation zwischen Fahrzeugen," *it—Information Technology (formerly it + ti)—Methoden und innovative Anwendungen der Informatik und Informationstechnik*, vol. 44, pp. 278–286, Oct. 2002.

[C27] F. Sivrikaya and B. Yener, "Time synchronization in sensor networks: A survey," *IEEE Network*, vol. 18, pp. 45–50, July–Aug. 2004.

Data Collection, Routing, and Architectures

[C28] R.C. Shah, S. Roy, S. Jain, and W. Brunette, "Data MULEs: Modeling and analysis of a three-tier architecture for sparse sensor networks," in *Ad Hoc Networks Journal*, vol. 1, pp. 215–233, Elsevier, Sept. 2003.

[C29] A. Chakrabarti, A. Sabharwal, and B. Aazhang, "Using predictable observer mobility for power efficient design of sensor networks," in *Information Processing in Sensor Networks (IPSN'03)*, Palo Alto, CA, Apr. 2003.

[C30] O. Younis and S. Fahmy, "HEED: A Hybrid, Energy-Efficient, Distributed Clustering Approach for Ad-hoc Sensor Networks," in *IEEE Transactions on Mobile Computing*, vol. 3, pp. 366–379, 2004.

[C31] R. Govindan, J. Hellerstein, W. Hong, S. Madden, M. Franklin, and S. Shenker, "The Sensor Network as a Database," Tech. Rep. 02–771, University of Southern California, 2002. ftp://ftp.usc.edu/pub/csinfo/tech-reports/papers/02-771.pdf.

[C32] M. Yarvis and W. Ye, "Tiered Architectures in Sensor Networks," in *Handbook of Sensor Networks: Compact Wireless and Wired Sensing Systems*, M. Ilyas and I. Mahgoub, eds., Boca Raton, FL, pp. 13.1–13.22, CRC Press, 2004.

Security

[C33] F. Stajano and R. Anderson, "The resurrecting duckling: Security issues for ad-hoc wireless networks," in *7th International Workshop on Security Protocols*, Cambridge, UK, Apr. 1999.

[C34] T. Martin, M. Hsiao, D. Ha, and J. Krishnaswami, "Denial-of-service attacks on battery-powered mobile computers," in *Proceedings of the 2nd IEEE Pervasive Computing Conference*, Orlando, FL, pp. 309–318, Mar. 2004.

Hardware

[D35] C. Schurgers, O. Aberthorne, and M. Srivastava, "Modulation scaling for energy aware communication systems," in *Proceedings of the 2001 International Symposium on Low Power Electronics and Design*, Huntington Beach, CA, pp. 96–99, Aug. 2001.

[D36] A.P. Chandrakasan, R. Min, M. Bhardwaj, S. Cho, and A. Wang, “Power aware wireless microsensor systems,” in *28th European Solid-State Circuits Conference (ESSCIRC’02)*, Florence, Italy, 2002.

[D37] S. Roundy, P. Wright, and J. Rabaey, “A study of low level vibrations as a power source for wireless sensor nodes,” *Computer Communications*, vol. 26, pp. 1131–1144, July 2003.

[D38] J. Kymissis, C. Kendall, J. Paradiso, and N. Gershenfeld, “Parasitic power harvesting in shoes,” in *Proceedings of the 2nd IEEE International Symposium on Wearable Computers (ISWC’04)*, Pittsburgh, PA, Oct. 1998.

[D39] A.J. Appleby, *Fuel Cell Handbook*, New York, NY: Van Reinhold Co., 1989.

[D40] W.C. Chung and D.S. Ha, “An Accurate Ultra WideBand (UWB) Ranging for precision asset location,” in *International Conference on UWB Systems and Technologies*, Reston, VA, Nov. 2002.

[D41] M. Vieira, D. da Silva Jr., C.C. Jr., and J. da Mata, “Survey on wireless sensor network devices,” in *Proceedings of the 9th IEEE International Conference on Emerging Technologies and Factory Automation (ETFA’03)*, Lisbon, Portugal, Sept. 2003.

[D42] B.A. Warneke and K.S.J. Pister, “MEMS for distributed wireless sensor networks,” in *Proceedings of the 9th International Conference on Electronics, Circuits and Systems (ICECS’02)*, vol. 1, Dubrovnik, Croatia, pp. 291–294, 2002.

[D43] Z. Karakehayov, “Low-Power Design for Smart Dust Networks,” in *Handbook of Sensor Networks: Compact Wireless and Wired Sensing Systems*, M. Ilyas and I. Mahgoub, eds., Boca Raton, FL, pp. 37.1–37.12, CRC Press, 2004.

[D44] B. Warneke, “Miniaturizing Sensor Networks with MEMS,” in *Handbook of Sensor Networks: Compact Wireless and Wired Sensing Systems*, M. Ilyas and I. Mahgoub, eds., Boca Raton, FL, pp. 5.1–5.19, CRC Press, 2004.

[D45] R. Min, M. Bhardwaj, S. Cho, A. Sinha, E. Shih, A. Wang, and A.P. Chandrakasan, “An Architecture for a Power-Aware Distributed Microsensor Node,” in *IEEE Workshop on Signal Processing Systems (SiPS’00)*, Lafayette, LA, Oct. 2000.

[D46] D.D. Wentzloff, B.H. Calhoun, R. Min, A. Wang, N. Ickes, and A.P. Chandrakasan, “Design considerations for next generation wireless power-aware microsensor nodes,” in *Proceedings of the 17th International Conference on VLSI Design*, Mumbai, India, pp. 361–367, 2004.

Existing Platforms

[E47] J. Beutel, O. Kasten, M. Ringwald, F. Siegemund, and L. Thiele, “Poster abstract: Btnodes—a distributed platform for sensor nodes,” in *Proceedings of the First International Conference on Embedded Networked Sensor Systems (SenSys-03)*, Los Angeles, CA, Nov. 2003.

[E48] V. Handziski, J. Polastre, J.-H. Hauer, C. Sharp, A. Wolisz, and D. Culler, “Flexible hardware abstraction for wireless sensor networks,” in *Proceedings of the 2nd International Workshop on Wireless Sensor Networks (EWSN 2005)*, Istanbul, Turkey, Jan. 2005.

[E49] R.M. Kling, “Intel Mote: An Enhanced Sensor Network Node,” in *International Workshop on Advanced Sensors, Structural Health Monitoring and Smart Structures at Keio University*, Tokyo, Japan, Nov. 2003.

[E50] K. Dantu, M. Rahimi, H. Shah, S. Babel, A. Dhariwal, and G. Sukhatme, “Robomote: Enabling Mobility In Sensor Networks,” Tech. Rep. CRES-04-006, University of Southern California.

[E51] M.B. McMickell, B. Goodwine, and L.A. Montestruque, “MICAbot: A robotic platform for large-scale distributed robotics,” in *Proceedings of International Conference on Intelligent Robots and Systems (ICRA’03)*, vol. 2, Taipei, Taiwan, pp. 1600–1605, 2003.

[E52] S. Bergbreiter and K.S.J. Pister, “CotsBots: An Off-the-Shelf Platform for Distributed Robotics,” in *Proceedings of the 2003 IEEE International Conference on Intelligent Robots and Systems (ICRA’03)*, Las Vegas, NV, Oct. 2003.

[E53] A. Savvides and M.B. Srivastava, “A distributed computation platform for wireless embedded sensing,” in *20th International Conference on Computer Design (ICCD’02)*, Freiburg, Germany, Sept. 2002.

[E54] S. Park, I. Locher, and M. Srivastava, “Design of a wearable sensor badge for smart kindergarten,” in *6th International Symposium on Wearable Computers (ISWC2002)*, Seattle, WA, pp. 13.1–13.22, Oct. 2002.

[E55] L.F.W. van Hoesel, S.O. Dulman, P.J.M. Havinga, and H.J. Kip, “Design of a low-power testbed for Wireless Sensor Networks and verification,” Tech. Rep. R-CTIT-03-45, University of Twente, Sept. 2003.

[E56] M. Strohbach, “The smart-its platform for embedded context-aware systems,” in *Proceedings of the First International Workshop on Wearable and Implantable Body Sensor Networks*, London, UK, Apr. 2004.

[E57] A. Schmidt, M. Strohbach, K.V. Laerhoven, and H.-W. Gellersen, “Ubiquitous interaction—Using surfaces in everyday environments as pointing devices,” in *7th ERCIM Workshop “User Interfaces For All,”* Chantilly, France, 2002.

[E58] M. Beigl, A. Krohn, T. Zimmer, C. Decker, and P. Robinson, “AwareCon: Situation aware context communication,” in *The Fifth International Conference on Ubiquitous Computing (Ubicomp’03)*, Seattle, WA, Oct. 2003.

[E59] J. Beutel, O. Kasten, F. Mattern, K. Roemer, F. Siegemund, and L. Thiele, “Prototyping sensor network applications with BTnodes,” in *IEEE European Workshop on Wireless Sensor Networks (EWSN’04)*, Berlin, Germany, Jan. 2004.

[E60] J. Lifton, D. Seetharam, M. Broxton, and J. Paradiso, “Pushpin computing system overview: A platform for distributed, embedded, ubiquitous sensor networks,” in *Proceedings of the Pervasive Computing Conference*, Zurich, Switzerland, Aug. 2002.

[E61] J.A. Paradiso, J. Lifton, and M. Broxton, “Sensate media—multimodal electronic skins as dense sensor networks,” *BT Technology Journal*, vol. 22, pp. 32–44, Oct. 2004.

[E62] K.V. Laerhoven, N. Villar, and H.-W. Gellersen, “Pin&Mix: When Pins Become Interaction Components. . . ,” in *Physical Interaction (PI03)—Workshop on Real World User Interfaces”—Mobile HCI Conference*, Udine, Italy, Sept. 2003.

Daniele Puccinelli received a Laurea degree in Electrical Engineering from the University of Pisa, Italy, in 2001. After spending two years in industry, he joined the graduate program in Electrical Engineering at the University of Notre Dame, and received an M.S. Degree in 2005. He is currently working toward his Ph.D. degree. His research has focused on cross-layer approaches to wireless sensor network protocol design, with an emphasis on the interaction between the physical and the network layer.

Martin Haenggi received the Dipl. Ing. (M.Sc.) degree in electrical engineering from the Swiss Federal Institute of Technology in Zurich (ETHZ) in 1995. In 1995, he joined the Signal and Information Processing Laboratory at ETHZ as a Teaching and Research Assistant. In 1996 he earned the Dipl. NDS ETH (post-diploma) degree in information technology, and in 1999, he completed his Ph.D. thesis on the analysis, design, and optimization of cellular neural networks. After a postdoctoral year at the Electronics Research Laboratory at the University of California in Berkeley, he joined the Department of Electrical Engineering at the University of Notre Dame as an assistant professor in January 2001. For both his M.Sc.and his Ph.D. theses, he was awarded the ETH medal, and he received an NSF CAREER award in 2005. For 2005/06, he is a CAS Distinguished Lecturer. His scientific interests include networking and wireless communications, with an emphasis on ad hoc and sensor networks.