

WIRELESS INTEGRATED NETWORK SENSORS

For pervasive computing performance, exploit the physical limits of these densely distributed networks of embedded sensors, controls, and processors.

Wireless integrated network sensors (WINS) provide distributed network and Internet access to sensors, controls, and processors deeply embedded in equipment, facilities, and the environment. The WINS network represents a new monitoring and control capability for applications in such industries as transportation, manufacturing, health care, environmental oversight, and safety and security. WINS combine microsensor technology and low-power signal processing, computation, and low-cost wireless networking in a compact system.

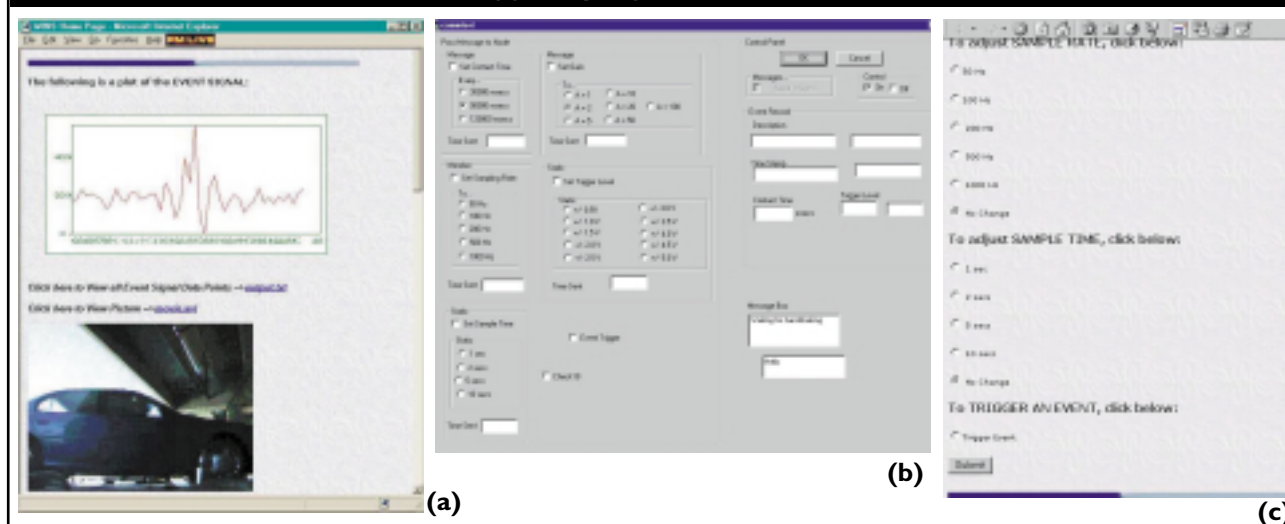
Recent advances in integrated circuit technology have enabled construction of far more capable yet inexpensive sensors, radios, and processors, allowing mass production of sophisticated systems linking the physical world to digital data networks

[2–5]. Scales range from local to global for applications in medicine, security, factory automation, environmental monitoring, and condition-based maintenance. Compact geometry and low cost allow WINS to be embedded and distributed at a fraction of the cost of conventional wireline sensor and actuator systems.

WINS opportunities depend on development of a scalable, low-cost, sensor-network architecture. Such applications require delivery of sensor information to the user at a low bit rate through low-power transceivers. Continuous sensor signal processing enables the constant monitoring of events in an environment in which short message packets would suffice. Future applications of distributed embedded processors and sensors will require vast numbers of devices. Conventional methods of sensor networking represent an impractical demand on cable installation and network bandwidth. Processing at the source would drastically reduce the financial, computational, and management burden on communication system

 G.J. Pottie and W.J. Kaiser

Figure 1. Remote access to seismic and image data;
(a) browser image of photo and seismic response of vehicle; **(b) sensor node control panel;** and
(c) sampling rate controls.



components, networks, and human resources.

Here, we limit ourselves to a security application designed to detect and identify threats within some geographic region and report the decisions concerning the presence and nature of such threats to a remote observer via the Internet. In the context of this application, we describe the physical principles leading to consideration of dense sensor networks, outline how energy and bandwidth constraints compel a distributed and layered signal processing architecture, outline why network self-organization and reconfiguration are essential, discuss how to embed WINS nodes in the Internet, and describe a prototype platform enabling these functions, including remote Internet control and analysis of sensor-network operation.

Physical Principles

When are distributed sensors better than a single large device, given the high cost of design implicit in having to create a self-organizing cooperative network? What are the fundamental limits in sensing, detection theory, communications, and signal processing driving the design of a network of distributed sensors?

Propagation laws for sensing. All signals decay with distance as a wavefront expands. For example, in free space, electromagnetic waves decay in intensity as the square of the distance; in other media, they are subject to absorption and scattering effects that can induce even steeper declines in intensity with distance. Many media are also dispersive (such as via multipath or low-pass filtering effects), so a distant sensor requires such costly operations as deconvolution (channel estimation and inversion) to partially

undo the dispersion [12]. Finally, many obstructions can render electromagnetic sensors useless. Regardless of the size of the sensor array, objects behind walls or under dense foliage cannot be detected.

As a simple example, consider the number of pixels needed to cover a particular area at a specified resolution. The geometry of similar triangles reveals that the same number of pixels is needed whether the pixels are concentrated in one large array or distributed among many devices. For free space with no obstructions, we would typically favor the large array, since there are no communications costs for moving information from the pixels to the processor. However, coverage of a large area implies the need to track multiple targets (a very difficult problem), and almost every security scenario of interest involves heavily cluttered environments complicated by obstructed lines of sight. Thus, if the system is to detect objects reliably, it has to be distributed, whatever the networking cost.

There are also example situations (such as radar) in which it is better to concentrate the elements, typically where it is not possible to get sensors close to targets. There are also many situations in which it is possible to place sensors in proximity to targets, bringing many advantages.

Detection and estimation theory fundamentals. A detector is given a set of observables $\{X_i\}$ to determine which of several hypotheses $\{h_i\}$ is true. These observables may, for example, be the sampled output of a seismic sensor. The signal includes not only the response to the desired target (such as a nearby pedestrian) but background noise and interference from other seismic sources. A hypothesis might include the

intersection of several distinct events (such as the presence of multiple targets of particular types).

The decision concerning target presence, absence, and type is usually based on estimates of parameters of these observations. Examples of parameters include selected Fourier, linear predictive coding, and wavelet transform coefficients. The number of parameters is typically a small fraction of the size of the observable set and thus constitute a reduced representation of the observations for purposes of distinguishing among hypotheses.

The set of parameters is known collectively as the feature set $\{f_k\}$. The reliability of this parameter estimation depends on both the number of independent observations and the signal-to-noise ratio (SNR). For example, according to the Cramer-Rao bound [10], which establishes the fundamental limits of estimation accuracy, the variance of a parameter estimate for a signal perturbed by white noise declines linearly with both the number of observations and the SNR. Consequently, to have to compute a good estimate of any particular feature, we need either a long set of independent observations or high SNR.

The formal means of choosing among hypotheses is to construct a decision space (whose coordinates are the values of the features) and divide it into regions according to the rule we decide on the hypothesis h_i , if the conditional probability $p(h_i|\{f_k\}) > p(h_j|\{f_k\})$ for all j not equal to i . Note that the features include environmental variations and other factors we measure or about which we have prior knowledge. The complexity of the decision increases with the dimension of the feature space; our uncertainty in the decision also generally increases with the number of hypotheses we have to sort through. Thus, to reliably distinguish among many possible hypotheses, we need a larger feature space. To build the minimum size space, we must determine the marginal improvement in the decision error rate resulting from addition of another feature. This may be as simple as including another term in an orthonormal expansion (such as fast Fourier transform and wavelet transform) or an entirely different transformation of the set $\{X\}$.

Unfortunately, we seldom know the prior probabilities of the various hypotheses; training is often inadequate to determine the conditional probabilities; and the marginal improvement in reliability declines rapidly as more features are extracted from any given set of observables.

On these facts hang many practical algorithms. For example, we could apply the deconvolution and target-separation machinery to exploit a distributed array. Though this machinery requires intensive communications and computations, it vastly reduces the

size of the feature space and the number of hypotheses that have to be considered, as each feature extractor deals with only one target with no propagation dispersal effects.

Alternatively, we may deploy a dense sensor network. Due to the decay of signals with distance, shorter-range phenomena (such as magnetics) can be used, limiting the number of targets (and hence hypotheses) in view at any given time. At short range, the probability is enhanced that the environment is essentially homogeneous within the detection range, reducing the number of environmental features—and thus the size of the decision space. Finally, since higher SNR is obtained at short range, and we can use a variety of sensing modes that may be unavailable at distance, we are better able to choose a small feature set that distinguishes targets. With only one mode, we would need to go deep into that mode's feature set, getting lower marginal returns for each feature. Thus, having targets nearby offers many options for reducing the size of the decision space.

Communications constraints. Spatial separation is another important factor in the construction of communication networks. For low-lying antennas, intensity drops as the fourth power of distance due to partial cancellation by a ground-reflected ray [7, 9]. Propagation is influenced by surface roughness, the presence of reflecting and obstructing objects, and antenna elevation. The losses make long-range communication a power-hungry exercise; the combination of Maxwell's Laws (governing propagation of electromagnetic radiation) and Shannon's capacity theorem (establishing fundamental relationships among bandwidth, SNR, and bit rate) together dictate that there is a limit on how many bits can be conveyed reliably, given power and bandwidth restrictions. On the other hand, the strong decay of intensity with distance provides spatial isolation, allowing the reuse of frequencies throughout a network.

Multipath propagation (resulting from reflections off multiple objects) is a serious problem. A digital modulation requires a 40dB increase in SNR to maintain an error probability of 10^{-5} with Rayleigh distributed-amplitude fading of the signal due to multipath, compared to a channel with the same average path loss perturbed only by Gaussian noise. It is possible to recover most of this loss by means of "diversity" obtainable in any of the domains of space, frequency, and time, since, with sufficient separation, the multipath fade levels are independent. By spreading the information, the multiple versions experience different fading, so the result is more akin to the average. If nothing is done, the worst-case conditions dominate error probabilities.

For static sensor nodes, time diversity is not an option with respect to path losses, although it may be a factor in jamming and other types of interference. Likewise, spatial diversity is difficult to obtain, since multiple antennas are unlikely to be mounted on small platforms. Thus, diversity is most likely to be achieved in the frequency domain by, say, employing some combination of frequency-hopped spread spectrum, interleaving, and channel coding. Measures known to be effective against deliberate jamming are also generally effective against multipath fading and multiuser interference. This interference reflects the common problem of intermittent events of poor SNR.

“Shadowing,” or wavefront obstruction and confinement, and path loss can be dealt with by employing a multihop network. If nodes are placed randomly in an environment, some links to near neighbors are obstructed, while others present a clear line of sight. The greater the density, the closer the nodes and the

switch at high frequency and with large voltage swings dominate the chip’s power cost.

While Moore’s Law implies that transistor areas continue to decrease—and signal-processing power costs decline with time—for radios and any communication technology, there are limits on the power required to transmit reliably over any given distance. The power-amplifier stage cannot be made smaller, due to limits on the current density of semiconductors. This stage typically burns at least four times the radiated energy, and so, in time, dominates the energy cost of radios. However, if we consider short-range communication with peak radiated power of less than 1mW, we continue to find that the oscillators and mixers used for up and down conversion dominate the energy budget; radios consume essentially the same power whether transmitting or receiving. While radio efficiency improves over time, with continuing technological advances, these facts suggest that net-

> If the system is to detect objects reliably, it has to be distributed, whatever the networking cost.

greater the likelihood of having a link with sufficiently small distance and shadowing losses. The signals then effectively hop around obstacles. Exploitation of these forms of diversity can lead to orders of magnitude reduction in the energy required to transmit data from one location in a WINS network to another; the energy cost is then dominated by the reception and retransmission energy costs of the radio transceivers for dense peer-to-peer networks.

Radio systems involve a close connection between networking strategy and physical layer. The connection is even stronger in light of the multiple-access nature of the channel, since interference among users is often the limiting impairment; the management of multiple-access interference is explored in [6].

Energy consumption in integrated circuits. Unfortunately, there are limits to the energy efficiency of complimentary metal oxide semiconductor (CMOS) communications and signal-processing circuits. Overall system costs cannot be low if the energy system is large. A CMOS transistor pair draws power each time it is flipped. The power used is roughly proportional to the product of the switching frequency, the area of the transistor (related to device capacitance), and the square of the voltage swing. Thus, power consumption for any given operation drops roughly as the fourth power of feature size. The components that

works should be designed so the radio is off as much of the time as possible and otherwise transmits only at the minimum required level.

Processing also gets cheaper with time but is not yet free. Because application-specific integrated circuits (ASICs) can clock at much lower speeds and use less numerical precision, they consume several orders of magnitude less energy than digital signal processors (DSPs). While the line between dedicated processors and general-purpose (more easily programmed) machines is constantly shifting, generally speaking, a mixed architecture is needed for computational systems dealing with connections to the physical world. The ratio in die area between the two approaches—ASIC and DSP—scales with technological change, so ASICs maintain a cost advantage over many chip generations. Convenient programmability across several orders of magnitude of energy consumption and data processing requirements is a worthy research goal for pervasive computing. In the meantime, while researchers continue to pursue that goal, multiprocessor systems are needed in WINS.

Signal-Processing Architectures

Security applications require constant vigilance by at least a subset of the sensors. We want a low false-alarm rate and a high detection probability. So as

long as data is queued, we can usually run energy-efficient procedures providing high detection probabilities—and high false-alarm rates. Energy thresholding and limited frequency analysis on low-sampling-rate magnetic, acoustic, infrared, and seismic sensors are excellent candidates for low-power ASICs. Higher-energy processing and sensing can be invoked if certainty levels are not high enough. Next, a WINS node might seek information from nearby sensors for data fusion (the weighted merging of detection decisions) or coherent beamforming (the complex weighting of raw data from multiple sensors for improved detection and target location). This cooperative behavior is a later step, since communication of raw data is very costly in terms of energy, as is its processing. Finally, a classification decision might be made using a large neural network or some other sophisticated procedure to provide the required degree of certainty. In the worst case, raw data may be hopped back to a remote site where a human would perform the pattern recognition. We can stop at any point in this chain when certainty thresholds are met.

Two design principles emerge from this effort to achieve reliable decisions with low energy consumption. First, we should play the probability game only to the extent we have to. Most of the time, there are no targets and thus no need to apply our most expensive algorithm (data to humans), but there are too many circumstances in which the least-expensive algorithm would fail. A processing hierarchy can lead to huge cost reductions while assuring the required level of reliability. Second, the processing hierarchy is intertwined with networking and data storage issues. For how long and at what location data is queued depends on where in the processing hierarchy the operation resides, whether a node communicates and to which set of neighboring nodes depends on the signal processing task. The communication costs in turn influence the processing strategy, including our willingness to communicate and whether the processing is centralized or distributed. Optimization is mandated by the physical constraints of the WINS network. Therefore the physical layer intrudes up through the network and signal-processing layers to applications.

To make concrete the effect of these constraints, assume the following: a 1GHz carrier frequency; an antenna elevation of $1/2$ wavelength; an efficient digital modulation, such as binary-phase-shift-keying (BPSK) transmission, 10^{-6} error probability, fourth-power distance loss, Rayleigh fading, and an ideal (no-noise) receiver. The energy cost of transmitting 1Kb a distance of 100 meters is approximately 3 joules. By contrast, a general-purpose processor with

100MIPS/W power could efficiently execute 3 million instructions for the same amount of energy. If the application and infrastructure permit, it pays to process the data locally to reduce traffic volume and make use of multihop routing and advanced communications techniques, such as coding, to reduce energy costs.

Indeed, exploitation of the application makes possible low-power design. For example, consider the situation of a remote security operation (see Figure 1). The figure's screen images display remote WINS Internet operation. The browser screen images (a) and (b) display events captured by an intelligent WINS node. For this system, the WINS node carries two sensors with seismic and imaging capability. The basic idea is that the seismic sensor is constantly vigilant, as it requires little power. Simple energy detection can be used to trigger the camera's operation. The image and the seismic record surrounding the event can then be communicated to a remote observer. In this way, the remote node needs to perform simple processing at low power, and the radio does not need to support the continuous transmission of images. The networking allows human (or computer) observers to be remote from both the scene and the storage of archival records. The image data allows verification of events and is usually required in security applications involving human response.

The seismic record and image of a vehicle (a) and a running human (b) creating the record are both shown in the figure. The WINS node and WINS-gateway node control Web pages (c) and (d), allowing direct and remote control (via the WINS network and the Internet) of event-recognition algorithms. For example, the seismic energy threshold for triggering an image can be controlled remotely. The number of images transmitted can be reduced with an increased sensor suite of short-range detectors, including infrared and magnetic, and by adding more sophisticated processing within the nodes.

Collaborative processing can extend the effective range of sensors and enable new functions. For example, consider the problem of target location. With a dense array, target position can be tracked by having all nodes detecting a disturbance make a report. The centroid of all nodes reporting the target is one possible estimate of the target's position. This detection technique requires the exchange of very few bits of information per node. Much more precise position estimates can be achieved through beamforming, which requires the exchange of time-stamped raw data among the nodes. Although the related processing is also much more costly, it yields higher SNR-processed data for subsequent classification decisions, long-range position location, and even some self-location

and calibration options for the nodes [11].

Depending on the application, it might be better to have sparse clusters of beamforming-capable nodes, rather than a dense deployment of less-intelligent nodes, or it may be better still to enable both sets of functions simultaneously. For example, we may overlay a less dense array of intelligent nodes commanding the capture of coherent data for purposes of beamforming. Allowing for heterogeneity from the outset greatly expands the processing horizons.

WINS Network Architecture

Unlike conventional wireless networks, a WINS network has to support large numbers of sensors in a local area with short range and low average bit-rate communication (fewer than 1–100Kbps). The network design has to address the requirement of servicing dense sensor distributions, emphasizing recovery of environmental information. In WINS networks, as a rule, we seek to exploit the short-distance separation between nodes to provide multihop communication through the power advantages outlined earlier. Since for short hops, transceiver power

one low-power protocol suite embodying these principles has been developed, including boot-up, MAC, energy-aware routing, and interaction with mobile units [8]. Its development indicates the feasibility of achieving distributed low-power operation in a flat multihop network.

Also clear is that for a wide range of applications, some way has to be found to conveniently link sensor networks to the Internet. Inevitably, some layering of the protocols (and devices) is needed to make use of these standard interfaces. For example, the WINS NG (next-generation) node architecture design (discussed later) addresses the constraints on robust operation, dense and deep distribution, interoperability with conventional networks and databases, operating power, scalability, and cost (see Figure 2). WINS gateways provide support for the WINS network and access between conventional network physical layers and their protocols and between the WINS physical layer and its low-power protocols. WINS system design exploits the reduced link range available through multihopping to provide advantages the system architect can choose from the following set:

> If the application and infrastructure permit, it pays to process the data locally to reduce traffic volume and make use of multihop routing and advanced communication techniques to reduce energy costs.

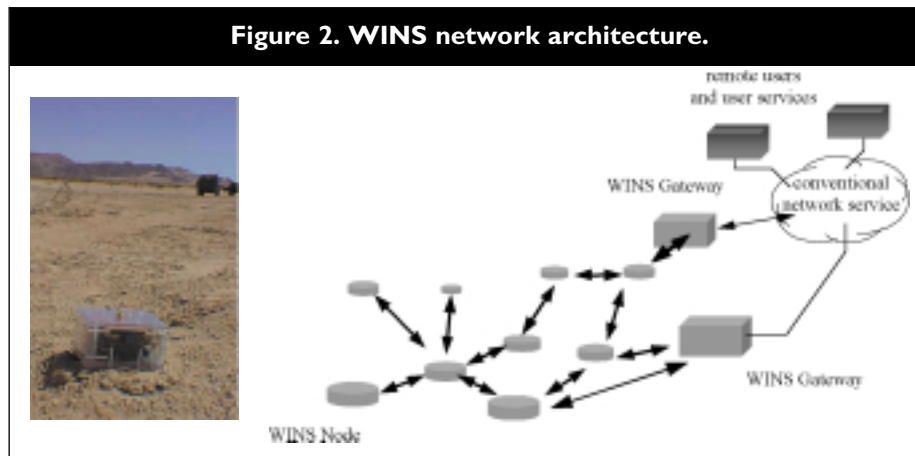
consumption for reception is nearly the same as that of transmission, the protocol should be designed so radios are off as much of the time as possible. That is, a device's medium access control (MAC) address in a network should include some variant of time-division multiple access.

A time-division protocol requires that the radios exchange short messages periodically to maintain local synchronism. It is not necessary for all nodes to have the same global clock, but the local variations from link to link should be small to minimize the guard times between slots and enable cooperative signal processing functions, including fusion and beamforming. The messages can combine network performance information, maintenance of synchronization, and reservation requests for bandwidth for longer packets. The abundant bandwidth resulting from the spatial reuse of frequencies and local processing ensures relatively few conflicts in these requests, so simple mechanisms can be used. At least

reduced operating power, improved bit rate, improved bit error rate, improved communication privacy (by way of reduced transmit power), simplified protocols, and reduced cost. These benefits are not obtained simultaneously but need to be extracted individually, depending on design emphasis.

In network design today, architects also have to address: How can Internet protocols, including TCP and IPv6, be employed within sensor networks? While it is desirable to not have to develop new protocols or perform protocol conversion at gateways, several factors demand custom solutions. First, IPv6 is not truly self-assembling; while addresses can be obtained from a server, this particular protocol presupposes attachment at lower levels already. Second, present-day Internet protocols take little account of the unreliability of physical channels or the need to conserve energy, focusing instead on support for a wide range of traffic. Embedded systems can achieve far higher efficiencies by exploiting the traffic's limited nature.

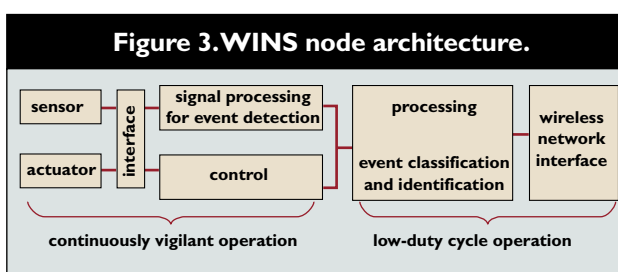
Another question they have to address is: Where should the processing and storage take place? Communication costs a great deal compared to processing; therefore energy constraints dictate doing as much processing at the source as possible. Moreover, reducing the amount of data to transmit simplifies network design significantly, permitting scalability to thousands of nodes per Internet gateway.



WINS Node Architectures

WINS development was initiated in 1993 at the University of California, Los Angeles; the first generation of field-ready WINS devices and software was fielded there three years later (see Figure 2a). The DARPA-sponsored low-power wireless integrated microsensors (LWIM) project demonstrated the feasibility of multihop, self-assembled, wireless networks. This first network also demonstrated the feasibility of algorithms for operating wireless sensor nodes and networks at micropower levels. In another DARPA-funded joint development program (involving UCLA and the Rockwell Science Center of Thousand Oaks, Calif.), a modular development platform was devised to enable evaluation of more sophisticated networking and signal-processing algorithms and to deal with many types of sensors, though with less emphasis on power conservation than LWIM [1]. These experiments taught us to recognize the importance of separating the real-time functions that have to be optimized for low power from the higher-level functions requiring extensive software development but that are invoked with light-duty cycles.

The WINS NG node architecture was subsequently developed by Sensor.com, founded by the authors in 1998 in Los Angeles, to enable continuous sensing, signal processing for event detection, local control of actuators, event identification, and communication at low power (see Figure 3). Since the event-detection process is continuous, the sensor, data converter, data buffer, and signal processing all have to operate at micropower levels, using a real-time system. If an event is detected, a process may be alerted to identify the event. Protocols for node operation then determine whether extra energy should be expended for further processing and whether a remote user or neighboring WINS node should be alerted. The



WINS node then communicates an attribute of the identified event, possibly the address of the event in an event look-up table stored in all network nodes.

These infrequent events can be managed by the higher-level processor—in the first version of WINS NG, a Windows CE-based device selected for the availability of low-cost developer tools. By providing application programming interfaces enabling the viewing and controlling of the lower-level functions, a developer is either shielded from real-time functions or is allowed to delve into them as desired to improve an application's efficiency. Future generations will also support plug-in Linux devices; other development will include very small but limited sensing devices that interact with WINS NG nodes in heterogeneous networks, supporting, say, intelligent tags (see Borriello's and Want's "Embedded Computation Meets the World Wide Web" in this issue). These small devices might scavenge their energy from the environment by means of photocells or piezoelectric materials, capturing energy from vibrations and achieving perpetual life spans. A clear technical path exists today, offering increased circuit integration and improved packaging. This path should produce very low-cost and compact devices in the near future.

Conclusion

These physical considerations are making it possible for us to pursue the innovative design of densely distributed sensor networks and the resulting advantages

of layered and heterogeneous processing and networking architectures for related applications. The close intertwining of network processing is a central feature of systems connecting the physical and virtual worlds. Development platforms are now available that will increasingly allow a broader community to engage in fundamental research in networking and new applications, advancing developers and users alike toward truly pervasive computing. ■

REFERENCES

1. Agre, J., Clare, L., Pottie, G., and Romanov, N. Development platform for self-organizing wireless sensor networks. In *Proceedings of Aerosense'99* (Orlando Fla., Apr. Apr. 8–9). International Society of Optical Engineering, Bellingham, Wa., 1999, 257–268.
2. Asada, G., Dong, M., Lin, T., Newberg, F., Pottie, G., Marcy, H., and Kaiser, W. Wireless integrated network sensors: Low-power systems on a chip. In *Proceedings of the 24th IEEE European Solid-State Circuits Conference* (Den Hague, The Netherlands, Sept. 21–25). Elsevier, 1998, 9–12.
3. Bult, K., Burstein, A., Chang, D., Dong, M., Fielding, M., Kruglick, E., Ho, J., Lin, F., Lin, T.-H., Kaiser, W., Marcy, H., Mukai, R., Nelson, P., Newberg, F., Pister, K., Pottie, G., Sanchez, H., Stafsudd, O., Tan, K., Ward, C., Xue, S., and Yao, J. Low-power systems for wireless microsensors. In *Proceedings of the International Symposium on Low-Power Electronics and Design* (Monterey, Calif., Aug. 12–14). IEEE, New York, 1996, 17–21.
4. Dong, M., Yung, G., and Kaiser, W. Low-power signal processing architectures for network microsensors. In *Proceedings of the 1997 International Symposium on Low-Power Electronics and Design* (Monterey, Calif., Aug. 18–20). IEEE, New York, 1997, 173–177.
5. Lin, T.-H., Sanchez, H., Rofougaran, R., and Kaiser, W. CMOS front-end components for micropower RF wireless systems. In *Proceedings of the 1998 International Symposium on Low-Power Electronics and Design*

- (Monterey, Calif., Aug. 10–12). IEEE, New York, 1998, 11–15.
6. Pottie, G. Wireless multiple access adaptive communication techniques. In *Encyclopedia of Telecommunications*, F. Froelich and A. Kent Eds. Marcel Dekker, Inc., New York, 1999, 1–41.
7. Rappaport, T. *Wireless Communications: Principles and Practice*. Prentice Hall, Upper Saddle River, N.J., 1996.
8. Sohrabi, K., Gao, J., Ailawadhi, V., and Pottie, G. A self-organizing sensor network. In *Proceedings of the 37th Allerton Conference on Communication, Control, and Computing* (Monticello, Ill., Sept. 27–29). Dept. of Electrical and Computer Engineering, University of Illinois at Urbana-Champaign, 1999.
9. Sohrabi, K., Manriquez, B., and Pottie, G. Near-ground wideband channel measurements. In *Proceedings of the 49th Vehicular Technology Conference* (Houston, May 16–20). IEEE, New York, 1999, 571–574.
10. Van Trees, H. *Detection, Estimation and Modulation Theory*. John Wiley & Sons, Inc., New York, 1968.
11. Yao, K., Hudson, R., Reed, C., Chen, D., and Lorenzelli, F. Blind beamforming on a randomly distributed sensor array system. *IEEE J. Select. Areas Comm.* 16, 8 (Oct. 1998), 1555–1567.
12. Yu, T., Chen, D., Pottie, G., and Yao, K. Blind decorrelation and deconvolution algorithm for multiple-input, multiple-output system. In *Proceed. Internat. Soc. Opt. Eng.* 3807 (July 1999), 200–209.

G.J. POTTIE (pottie@ee.ucla.edu) is an associate professor in the Electrical Engineering Department at the University of California, Los Angeles, and a founder of Sensoria Corp. in Los Angeles.

W.J. KAISER (kaiser@ee.ucla.edu) is a professor in and department chair of the Electrical Engineering Department at the University of California, Los Angeles, and a founder of Sensoria Corp. in Los Angeles.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

© 2000 ACM 0002-0782/00/0500 \$5.00

The Computer Image

Alan Watt, Fabio Polcarpo


To Order

Call Toll Free:

1.800.342.6626 (USA & Canada)

1.212.626.0500 (outside US)

Fax: 1.212.944.1318



1998

751 pages Hardcover

ISBN 0-201-42298-0

ACM Order Number: 702980

Member Price: \$53.95

Nonmember Price: \$59.95

To order online: <http://www.acm.org/catalog/>

ad2

Association for Computing Machinery
The First Society in Computing

