INSTRUMENTING THE WORLD WITH WIRELESS SENSOR NETWORKS

*D. Estrin, L. Girod**

UCLA
Department of Computer Science
email: {*destrin,girod*}@*cs.ucla.edu*

G. Pottie, M. Srivastava†

UCLA
Department of Electrical Engineering
email: {*mbs,pottie*}@*ee.ucla.edu*

ABSTRACT

Pervasive micro-sensing and actuation may revolutionize the way in which we understand and manage complex physical systems: from airplane wings to complex ecosystems. The capabilities for detailed physical monitoring and manipulation offer enormous opportunities for almost every scientific discipline, and it will alter the feasible granularity of engineering.

We identify opportunities and challenges for distributed signal processing in networks of these sensing elements and investigate some of the architectural challenges posed by systems that are massively distributed, physically-coupled, wirelessly networked, and energy limited.

1. INTRODUCTION

The availability of low-power micro-sensors, actuators, embedded processors, and radios is enabling the application of distributed wireless sensing to a wide range of applications, including environmental monitoring, smart spaces, medical applications, and precision agriculture [1][2]. Most deployed sensor networks involve relatively small numbers of sensors, wired to a central processing unit where all of the signal processing is performed [3]. In contrast, this paper focuses on **distributed, wireless, sensor networks** in which the signal processing is **distributed** along with the sensing.

- **Why distributed sensing?** When the precise location of a signal of interest is unknown in a monitored region, distributed sensing allows one to place the sensors closer to the phenomena being monitored than if only a single sensor were used. This yields higher SNR, and improved opportunities for line of sight. While SNR can be addressed in many cases by deploying one very large sensitive sensor, line of sight, and more generally obstructions, cannot be addressed by deploying one sensor regardless of its sensitivity. Thus, distributed sensing provides robustness to environmental obstacles.
- **Why wireless?** When wired networking of distributed sensors can be easily achieved, it is often the more advantageous approach. Moreover, when nodes can be wired to renewable (relatively infinite) energy sources, this too greatly simplifies the system design and operation. However, in many envisioned applications, the environment being monitored does not have installed infrastructure for either communications or energy, and therefore untethered nodes must rely on local, finite, and relatively small energy sources, as well as wireless communication channels.
- **Why distributed processing?** Finally, although sensors are distributed to be close to the phenomena, one might still consider an architecture in which sensor outputs could be communicated back to a central processing unit. However, in the context of untethered nodes, the finite energy budget is a primary design constraint. Communications is a key energy consumer as the radio signal power in sensor networks drops off with r^4[4] due to ground reflections from short antenna heights. Therefore, one wants to process data as much as possible inside the network to reduce the number of bits transmitted, particularly over longer distances.

2. MOTIVATING APPLICATION

The potential applications of wireless sensor networks are highly varied: e.g., Physiological monitoring; Environmental monitoring (air, water, soil chemistry); Condition based maintenance; Smart spaces; Military; Precision agriculture; Transportation; Factory instrumentation and inventory tracking

Habitat monitoring [Cerpa-etal01, Hamilton, Steere-etal00] provides a rich collection of sensing modalities and environmental conditions and we use it to motivate our technical discussion. Consider the goal of supporting data collection and model development of complex ecosystems. Scientists and environmental impact monitoring authorities would like to monitor soil and air chemistry, as well as plant and animal species populations and behavior. For the latter, the primary modalities are imaging and acoustics to localize, identify and track species or phenomena based on implicit signals (acoustic and seismic), or explicit signals (RF tags). These facilities must be deployable in remote locations that lack installed energy and communication infrastructures, motivating the need for low-power wireless communication.

The strategy for node cooperation strategy has significant consequences in terms of communication bandwidth and energy consumption. For example, consider the task of identifying bird species in view of several cameras. If it is to be accomplished through image analysis, we could stream all the video back to a human operator-a very costly approach. Alternatively, we could stream audio to a central location, which then performs signal processing to identify and stream back only those streams that are most likely to contain a target species. While this reduces communications

*Supported by the DARPA SensIT program and NSF Special Projects under ANIR.

†Supported by the DARPA SensIT program.

overhead greatly, it still suffers from communications latency and lacks scalability due to the need to stream audio thorough a central processing point. Finally, we might distribute the problem further, hosting the audio signal processing software on the nodes, and developing algorithms that require only local cooperation to make a decision to capture images. This approach is scalable in that no long-range streaming of audio or video is necessary, resulting in more efficient use of communications bandwidth and limited energy resources.

In the remainder of this paper we identify some of the technical challenges associated with the design of wireless sensor networks and discuss several algorithmic approaches.

3. TECHNICAL CHALLENGES

Most envisioned sensor network applications encounter one or more of the following challenges:

- **Untethered** for energy and communication requiring maximal focus on energy efficiency.
- **Ad hoc** deployment, requiring that the system identifies and copes with the resulting distribution and connectivity of nodes.
- **Dynamic** environmental conditions requiring the system to adapt over time to changing connectivity and system stimuli.
- **Unattended** operation requiring configuration and reconfiguration be automatic (self-configuration)

To address these challenging environments, several strategies are likely to be key building blocks/techniques for wireless sensor networks:

- **Collaborative signal processing** among nodes that have experienced a common stimulus will greatly enhance the efficiency (information per bit transmitted) of these systems. Develop both **coherent** signal processing on small clusters by a centralized entity within the cluster, and **non-coherent** processing with much less stringent synchronization requirements and applicable across larger numbers of more loosely coupled elements.
- **Exploiting redundancy** of hardware elements to compensate for ad hoc deployment of systems. If elements cannot be carefully positioned relative to each other and the environment, then an alternate strategy to achieve "coverage" is to deploy a greater density of elements so that one can make use of some subset that have the desired absolute and relative position. In some contexts, even if elements can be uniformly placed in 3-space, environmental conditions might be such that coverage is not uniform due to obstacles and other sources of noise. Another application of redundancy is when the incremental cost of a node during initial deployment is much smaller than the incremental cost of deploying new nodes or renewing node resources (e.g., energy). In this case, one can exploit redundancy to extend system lifetime by adjusting duty cycle based on local density and local demand.
- **Adaptive fidelity signal processing** is another strategy that can be exploited in sensor networks to make trade-offs between energy, accuracy, and rapidity of results. Recognizing that one is trying to detect non-deterministic phenomena in the presence of communication noise and sensor diversity, the fidelity and timeliness of the signal processing at individual sensor nodes can be adapted to energy resources and latency requirements.
- **A hierarchical, tiered architecture** can greatly contribute to overall system lifetime and capability. Whenever possible, higher capacity system elements can be used to offload drain on small form factor elements, while the latter can be exploited to obtain the desired physical proximity to stimuli. Moreover, even among elements with homogenous capabilities, creating clusters and assigning special combining functions to cluster heads can contribute to overall system scalability. However, to avoid compromising robustness, such clusters/hierarchy must be self-configuring and reconfiguring in the face of environmental or network changes.

4. TECHNICAL APPROACHES

We now describe three generic techniques that would enable distributed signal processing tasks in wireless sensor networks.

4.1. Coherent processing algorithms

Coherent signal processing algorithms are distinguished from non-coherent methods in that information about the phase of the wavefront impinging on the nodes must be conveyed. Beamforming techniques allow localization of signals that originate within the convex hull of the participating nodes, higher SNR estimates of the signals compared to non-coherent methods, and determination of bearing angles for signals that originate outside the convex hull of the participating nodes. The price is a higher level of synchronization (to within a small fraction of one oscillation), and communication of relatively high bit rate data streams consisting of sampled waveforms. Given its high resource cost, we should resort to coherent processing only when we cannot attain adequate accuracy in the result with non-coherent methods such as combination of likelihood functions.

One way to organize the operations leading to coherent beamforming is as follows. Nodes go through a sequence of internal levels of signal processing before determining that neighbors should be involved in a detection/localization decision. An ad hoc network is constructed for non-coherent decision-making using for example the single winner election algorithm of [2]. The algorithm is optimized to minimize the overhead in finding a fusion center, since relatively little data must actually be communicated. However, if the decision has insufficient certainty or resolution, the same set of nodes become involved in a new network set-up that seeks to minimize the energy consumption in conveying sampled waveforms to a common central processing point. To this end significant overhead is acceptable since large amounts of data will be conveyed in the local neighborhood. A multi-winner election algorithm to accomplish this is also described in [2]. Standard beamforming techniques can now be applied using the data collected from the cluster of nodes.

There is no requirement for uniform lay-down of nodes to achieve beamforming [5]. To track distant sources, two or more clusters of nodes can be used, and with the intersection of the bearing lines used to establish location. Note that simply using all nodes in the network to do one massive beamforming operation could accomplish this end, but excessive communications and signal processing complexity would be required. Rather, for a scalable solution a signal processing step is required that recognizes whether near or far objects are being tracked. A crude technique is to consider the SNR variations among nodes in a cluster and to neighboring clusters. If the SNR is similar, then the signal source is likely to be distant. Having made this determination, clusters may decide to estimate lines of bearing or not, whether probabilistically or according to a predetermined schedule. The information on the bearing lines is then conveyed to a central node designated to perform the (noncoherent) fusion. Thus, there is never a case in which sampled waveforms must be conveyed over a large number of hops.

Achieving the required level of synchronism for coherent beamforming is in principle relatively straightforward for systems in which every node possesses a radio. Since the propagation velocity of seismic and acoustic signals is six orders of magnitude slower than that of radio waves, achieving data lock for RF communications would seem to already be much better synchronization than is required for beamforming. However, particular care must be paid to the node architecture to take advantage of this timing information. The typical interrupt cycles of general-purpose processors can be tens of milliseconds, an eternity with respect to even acoustic signals. Thus, embedded real-time components are required in the nodes to deal with time-stamping of the data.

4.2. Localization

Node location is employed by routing protocols that use spatial addresses, and by signal processing algorithms (e.g. beamforming) that are used for tasks such as target tracking. The underlying algorithm problem is that of localization whereby the nodes in the network discover their spatial coordinates upon network boot-up. When the sensor nodes are deployed in an unplanned topology, there is no a priori knowledge of location. The use of GPS in sensor nodes is ruled out in many scenarios because of power consumption, antenna size, and overhead obstructions such as dense foliage. The ad hoc nature of deployment rules out infrastructure for many scenarios of localization. It is critical that sensor network nodes be able to estimate their relative positions without assistance, using means that can be built-in.

The localization problem in itself is a good example of a signal-processing task that the sensor network needs to solve. The basic approach would be for sensor nodes to gather sufficient number of pair-wise distance estimates via some suitable mechanism, and then use multilateration algorithms to estimate positions of the nodes. To begin with, a few nodes might know their position via other means (beacon nodes), but at the end of the localization process every node would hopefully know its position.

A key problem however is that in conventional formulations of multilateration [6][7] one needs to estimate the location of an entity given estimates of its distance to 3 or more beacons with known positions. In sensor networks a very high density of beacons nodes would be needed. To keep the required beacon density and energies low, a preferred method would be to jointly estimate positions of all the non-beacon nodes via a collaborative multilateration formulation based on criterion such as least-square error minimization. Besides being computationally hard for large number of nodes, doing this would require a centralized node where all the distance estimates would be collected at significant communication and associated energy cost. A more scalable solution is locally distributed iterative multilateration [8] whereby a node calculates its position and is promoted to a beacon as soon as enough of its 1-hop neighbors are beacons. Starting with a critical density of beacons, a percolation-like phenomenon would result in gradually all the nodes discovering their position. With a sufficient beacon density, a small number of successive multilateration steps lead to rapid convergence of location estimates. The communication overhead is much lower than in centralized approach as all message exchange is strictly local and is easily piggybacked on routing messages.

Another challenge in localization is estimation of distance between a pair of nodes. Using time-of-flight of radio signals (as in GPS) is ruled out when the distances are too tiny and radio frequencies not very high. A readily available method would be to use the received signal strength indication (RSSI) provided be the radio. The RSSI data can be cheaply piggybacked on regular routing and data. The accuracy of this approach can be improved by using a parameterized channel, path loss model whose parameters are also estimated together with position [8]. However, in practice, the RSSI based approach works only in the absence of significant multipath effects. In most environments other than open spaces multipath is an issue. A promising alternative technology is to estimate distance by time of flight of acoustic or ultrasound signals, and using the much faster radio signal to establish time reference [9][10][11].

4.3. Distributed power management

Dynamic power management techniques such as shutdown and dynamic voltage scaling have emerged as powerful methods for power-aware computing. Power-aware operation is even more important for wireless sensor networks, and requires distributed versions of power management techniques.

As an example, consider shutdown, which is widely used in portable computing systems such as laptops. In sensor nets one could exploit redundant nodes by turning on only a time- varying subset of nodes, where the subset is selected for desired sensor and radio coverage. The remaining nodes can be shutdown, only to be woken up to provide additional sensor readings or communication routes when something interesting happens [12].

A key problem in such a distributed shutdown scheme is the strategy to select which node to shutdown and which to turn on at any given instant. A good way to model this problem is to optimally divide the sensor nodes into several subsets such that any given subset provides a baseline level of sensing and communication coverage. The different subsets can then be turned on and off according to a duty cycle determined by a repetitive schedule. As nodes die by depleting their batteries, the subsets are changed.

Unfortunately, modeling the problem in this fashion requires one to gather global information to find the subsets. Since communication is expensive in energy, the cost of the power management algorithm would swamp the savings from power management! **This illustrates the dilemma that so often arises in problems in sensor nets: the seemingly optimal way of solving a problem often results in algorithms whose communication energy costs exceed their benefits.** Therefore, a better strategy is to use algorithms that only shoot for good though sub optimal results but require only locally distributed processing with minimal communication costs.

This suggests that the decision regarding when to shutdown and wakeup a node should be made using information in the local neighborhood. A simple scheme is to turn the nodes on an off with a certain duty cycle with random phase differences. When off, the node power consumption can be reduced to microwatts. However, the power savings come at the cost of reduction in detection and classification accuracy as the sensor would miss a phenomenon when off. The cost of such missed events is very application specific: it may be okay to miss a frequent event, but very crucial to detect a rare event. In general, such a goal might be quantified as a quality of service requirement such as the following in the case of a sensor net for surveillance: "a target with 20 mph speed following this track will not pass undetected".

More sophisticated locally distributed schemes for shutdown would coordinate the on and off periods of neighboring nodes to improve energy efficiency for the same level of detection performance, and perhaps adapt the duty cycle parameters based on event activity. A crucial problem here is that of waking up the node. In shutdown in PCs and laptops external events such as keyboard presses or arrival of network packet result in the rest of the system waking up. However, in sensor nodes it is highly desirable to turn off the radio, which is usually more power-hungry than the processor and the sensors. Turning off the radio, unfortunately, means that a neighboring node that detected an interesting event cannot wake a node up. This can lead to missed events and packets. Therefore, one technological challenge for effective power management of sensor networks is to have an ultra low-power paging communication channel to wake up neighboring nodes on demand. An alternative is for low-power sensors to be constantly vigilant, with radio wake-up according to signal processing results.

The proposed form of adaptive duty cycle applies at multiple levels in the system. At short ranges, radios will consume nearly as much energy whether transmitting or receiving. Consequently, it makes sense to arrange for a time-division structure for communications. In [2] an algorithm is described for enabling distributed boot-up of the network, establishing both channel assignments and synchronism in an energy-efficient fashion. It assumes that sufficient signal processing is done in each node to result on average with traffic occupying a very small fraction of the available bandwidth. Nodes regularly communicate with their neighbors to keep synchronism and to indicate whether larger slots are to be reserved for bulk data transfers, in effect alerting the neighbor to turn on its receiver for some specified time period.

Maintaining a constant level of synchronism is helpful in reducing network latency and for health-keeping, but it is not the most energy efficient strategy in all traffic scenarios. When the signal sources of interest have a very low duty cycle, it may be better to periodically exchange packets to re-synch only to within some very coarse accuracy, and then spend significant energy in achieving tighter synchronism when there is something of interest to send [13].

5. CONCLUSIONS

In conclusion, wireless sensor networks present fascinating challenges for the application of distributed signal processing and distributed control. These systems will challenge us to apply appropriate techniques and metrics in light of the technology opportunities (cheap processing and sensing nodes) and challenges (energy constraints).

6. REFERENCES

[1] D. Estrin, R. Govindan, and J. Heidemann, Eds., *Special Issue on Embedding the Internet*. Communications of the ACM, vol. 43, no. 5, May 2000.

[2] B.R. Badrinath, M. Srivastava, K. Mills, J. Scholtz, and K. Sollins, Eds., *Special Issue on Smart Spaces and Environments*. IEEE Personal Communications, Oct. 2000.

[3] *Sensors: The Journal of Applied Sensing Technology*.

[4] G. Pottie and W. Kaiser, "Wireless sensor networks," *Communications of the ACM*, vol. 43, no. 5, pp. 51–58, May 2000.

[5] K. Yao, R. Hudson, C. Reed, D. Chen, and F. Lorenzelli, "Blind beamforming on a randomly distributed sensor array system," *IEEE JSAC*, vol. 16, no. 8, Oct. 1998.

[6] G.L. Stuber and J.J. Caffery jr., *The Mobile Communications Handbook, J.D. Gibson and E.M. Gibson, ed.*, chapter 24, *Radiolocation Techniques*, CRC Press, 1999.

[7] G. Turin, W. Jewell, and T. Johnston, "Simulation of urban vehicle-monitoring systems," *IEEE Transactions on Vehicular Technology*, vol. 21, no. 1, Feb. 1972.

[8] A. Savvides, F. Koushanfar, A. Boulis, V. Karavas, and M.B. Srivastava M. Potkonjak, "Location discovery in adhoc wireless networks," Memorandum, Networked and Embedded Systems Laboratory, UCLA, June 2000.

[9] A. Ward, A. Jones, and A. Hopper, "A new location technique for the active office," *IEEE Personal Communications*, vol. 4, no. 5, Oct. 1997.

[10] N. Priyantha, A. Chakraborty, and H. Balakrishnan, "The cricket location support system," in *Mobicom 2000*. ACM, Aug. 2000.

[11] L. Girod and D. Estrin, "Robust range estimation for localization in adhoc sensor networks," Tech. Rep. CS-TR-2000XX, UCLA, Nov. 2000.

[12] D. Estrin, R. Govindan, J. Heidemann, and S. Kumar, "Next century challenges: Scalable coordination in sensor networks," in *Mobicom 99*. ACM, Aug. 1999.

[13] J. Elson and D. Estrin, "Time synchronization for wireless sensor networks," in *IPDPS Workshop on Parallel and Distributed Computing Issues in Wireless Networks and Mobile Computing*, Apr. 2001.