1

Introduction

1.1 Wireless sensor networks: the vision

Recent technological advances allow us to envision a future where large numbers of low-power, inexpensive sensor devices are densely embedded in the physical environment, operating together in a wireless network. The envisioned applications of these wireless sensor networks range widely: ecological habitat monitoring, structure health monitoring, environmental contaminant detection, industrial process control, and military target tracking, among others.

A US National Research Council report titled *Embedded Everywhere* notes that the use of such networks throughout society "could well dwarf previous milestones in the information revolution" [47]. Wireless sensor networks provide bridges between the virtual world of information technology and the real physical world. They represent a fundamental paradigm shift from traditional inter-human personal communications to autonomous inter-device communications. They promise unprecedented new abilities to observe and understand large-scale, real-world phenomena at a fine spatio-temporal resolution. As a result, wireless sensor networks also have the potential to engender new breakthrough scientific advances.

While the notion of networking distributed sensors and their use in military and industrial applications dates back at least to the 1970s, the early systems were primarily wired and small in scale. It was only in the 1990s – when wireless technologies and low-power VLSI design became feasible – that researchers began envisioning and investigating large-scale embedded wireless sensor networks for dense sensing applications.

Figure 1.1 A Berkeley mote (MICAz MPR2400 series)

Perhaps one of the earliest research efforts in this direction was the low-power wireless integrated microsensors (LWIM) project at UCLA funded by DARPA [98]. The LWIM project focused on developing devices with low-power electronics in order to enable large, dense wireless sensor networks. This project was succeeded by the Wireless Integrated Networked Sensors (WINS) project a few years later, in which researchers at UCLA collaborated with Rockwell Science Center to develop some of the first wireless sensor devices. Other early projects in this area, starting around 1999–2000, were also primarily in academia, at several places including MIT, Berkeley, and USC.

Researchers at Berkeley developed embedded wireless sensor networking devices called motes, which were made publicly available commercially, along with TinyOS, an associated embedded operating system that facilitates the use of these devices [81]. Figure 1.1 shows an image of a Berkeley mote device. The availability of these devices as an easily programmable, fully functional, relatively inexpensive platform for experimentation, and real deployment has played a significant role in the ongoing wireless sensor networks revolution.

1.2 Networked wireless sensor devices

As shown in Figure 1.2, there are several key components that make up a typical wireless sensor network (WSN) device:

1. **Low-power embedded processor:** The computational tasks on a WSN device include the processing of both locally sensed information as well as information communicated by other sensors. At present, primarily due to economic

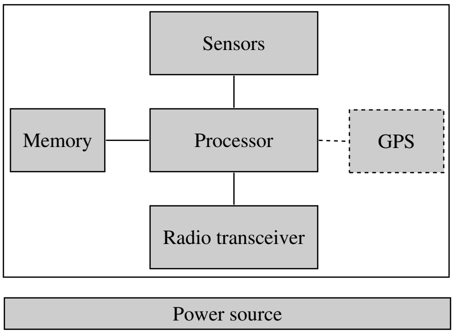

Figure 1.2 Schematic of a basic wireless sensor network device

constraints, the embedded processors are often significantly constrained in terms of computational power (e.g., many of the devices used currently in research and development have only an eight-bit 16-MHz processor). Due to the constraints of such processors, devices typically run specialized component-based embedded operating systems, such as TinyOS. However, it should be kept in mind that a sensor network may be heterogeneous and include at least some nodes with significantly greater computational power. Moreover, given Moore's law, future WSN devices may possess extremely powerful embedded processors. They will also incorporate advanced low-power design techniques, such as efficient sleep modes and dynamic voltage scaling to provide significant energy savings.

2. **Memory/storage:** Storage in the form of random access and read-only memory includes both program memory (from which instructions are executed by the processor), and data memory (for storing raw and processed sensor measurements and other local information). The quantities of memory and storage on board a WSN device are often limited primarily by economic considerations, and are also likely to improve over time.
3. **Radio transceiver:** WSN devices include a low-rate, short-range wireless radio (10–100 kbps, <100 m). While currently quite limited in capability too, these radios are likely to improve in sophistication over time – including improvements in cost, spectral efficiency, tunability, and immunity to noise, fading, and interference. Radio communication is often the most power-intensive operation in a WSN device, and hence the radio must incorporate energy-efficient sleep and wake-up modes.
4. **Sensors:** Due to bandwidth and power constraints, WSN devices primarily support only low-data-rate sensing. Many applications call for multi-modal sensing, so each device may have several sensors on board. The specific

sensors used are highly dependent on the application; for example, they may include temperature sensors, light sensors, humidity sensors, pressure sensors, accelerometers, magnetometers, chemical sensors, acoustic sensors, or even low-resolution imagers.

5. **Geopositioning system:** In many WSN applications, it is important for all sensor measurements to be location stamped. The simplest way to obtain positioning is to pre-configure sensor locations at deployment, but this may only be feasible in limited deployments. Particularly for outdoor operations, when the network is deployed in an *ad hoc* manner, such information is most easily obtained via satellite-based GPS. However, even in such applications, only a fraction of the nodes may be equipped with GPS capability, due to environmental and economic constraints. In this case, other nodes must obtain their locations indirectly through network localization algorithms.
6. **Power source:** For flexible deployment the WSN device is likely to be battery powered (e.g. using LiMH AA batteries). While some of the nodes may be wired to a continuous power source in some applications, and energy harvesting techniques may provide a degree of energy renewal in some cases, the finite battery energy is likely to be the most critical resource bottleneck in most WSN applications.

Depending on the application, WSN devices can be networked together in a number of ways. In basic data-gathering applications, for instance, there is a node referred to as the *sink* to which all data from *source* sensor nodes are directed. The simplest logical topology for communication of gathered data is a single-hop star topology, where all nodes send their data directly to the sink. In networks with lower transmit power settings or where nodes are deployed over a large area, a multi-hop tree structure may be used for data-gathering. In this case, some nodes may act both as sources themselves, as well as routers for other sources.

One interesting characteristic of wireless sensor networks is that they often allow for the possibility of intelligent in-network processing. Intermediate nodes along the path do not act merely as packet forwarders, but may also examine and process the content of the packets going through them. This is often done for the purpose of data compression or for signal processing to improve the quality of the collected information.

1.3 Applications of wireless sensor networks

The several envisioned applications of WSN are still very much under active research and development, in both academia and industry. We describe a few applications from different domains briefly to give a sense of the wide-ranging scope of this field:

1. **Ecological habitat monitoring:** Scientific studies of ecological habitats (animals, plants, micro-organisms) are traditionally conducted through hands-on field activities by the investigators. One serious concern in these studies is what is sometimes referred to as the "observer effect" – the very presence and potentially intrusive activities of the field investigators may affect the behavior of the organisms in the monitored habitat and thus bias the observed results. Unattended wireless sensor networks promise a cleaner, remote-observer approach to habitat monitoring. Further, sensor networks, due to their potentially large scale and high spatio-temporal density, can provide experimental data of an unprecedented richness.

   One of the earliest experimental deployments of wireless sensor networks was for habitat monitoring, on Great Duck Island, Maine [130]. A team of researchers from the Intel Research Lab at Berkeley, University of California at Berkeley, and the College of the Atlantic in Bar Harbor deployed wireless sensor nodes in and around burrows of Leach's storm petrel, a bird which forms a large colony on that island during the breeding season. The sensor-network-transmitted data were made available over the web, via a base station on the island connected to a satellite communication link.

2. **Military surveillance and target tracking:** As with many other information technologies, wireless sensor networks originated primarily in military-related research. Unattended sensor networks are envisioned as the key ingredient in moving towards network-centric warfare systems. They can be rapidly deployed for surveillance and used to provide battlefield intelligence regarding the location, numbers, movement, and identity of troops and vehicles, and for detection of chemical, biological, and nuclear weapons.

   Much of the impetus for the fast-growing research and development of wireless sensor networks has been provided though several programs funded by the US Defense Advanced Research Projects Agency (DARPA), most notably through a program known as Sensor Information Technology (SensIT) [188] from 1999 to 2002. Indeed, many of the leading US researchers and entrepreneurs in the area of wireless sensor networks today have been and are being funded by these DARPA programs.

3. **Structural and seismic monitoring:** Another class of applications for sensor networks pertains to monitoring the condition of civil structures [231]. The structures could be buildings, bridges, and roads; even aircraft. At present the health of such structures is monitored primarily through manual and visual

inspections or occasionally through expensive and time-consuming technologies, such as X-rays and ultrasound. Unattended networked sensing techniques can automate the process, providing rich and timely information about incipient cracks or about other structural damage. Researchers envision deploying these sensors densely on the structure – either literally embedded into the building material such as concrete, or on the surface. Such sensor networks have potential for monitoring the long-term wear of structures as well as their condition after destructive events, such as earthquakes or explosions. A particularly compelling futuristic vision for the use of sensor networks involves the development of controllable structures, which contain actuators that react to real-time sensor information to perform "echo-cancellation" on seismic waves so that the structure is unaffected by any external disturbance.

4. **Industrial and commercial networked sensing:** In industrial manufacturing facilities, sensors and actuators are used for process monitoring and control. For example, in a multi-stage chemical processing plant there may be sensors placed at different points in the process in order to monitor the temperature, chemical concentration, pressure, etc. The information from such real-time monitoring may be used to vary process controls, such as adjusting the amount of a particular ingredient or changing the heat settings. The key advantage of creating wireless networks of sensors in these environments is that they can significantly improve both the cost and the flexibility associated with installing, maintaining, and upgrading wired systems [131]. As an indication of the commercial promise of wireless embedded networks, it should be noted that there are already several companies developing and marketing these products, and there is a clear ongoing drive to develop related technology standards, such as the IEEE 802.15.4 standard [94], and collaborative industry efforts such as the Zigbee Alliance [244].

1.4 Key design challenges

Wireless sensor networks are interesting from an engineering perspective, because they present a number of serious challenges that cannot be adequately addressed by existing technologies:

1. **Extended lifetime:** As mentioned above, WSN nodes will generally be severely energy constrained due to the limitations of batteries. A typical alkaline battery, for example, provides about 50 watt-hours of energy; this may translate to less than a month of continuous operation for each node in full active mode. Given the expense and potential infeasibility of monitoring and

replacing batteries for a large network, much longer lifetimes are desired. In practice, it will be necessary in many applications to provide guarantees that a network of unattended wireless sensors can remain operational without any replacements for several years. Hardware improvements in battery design and energy harvesting techniques will offer only partial solutions. This is the reason that most protocol designs in wireless sensor networks are designed explicitly with energy efficiency as the primary goal. Naturally, this goal must be balanced against a number of other concerns.

2. **Responsiveness:** A simple solution to extending network lifetime is to operate the nodes in a duty-cycled manner with periodic switching between sleep and wake-up modes. While synchronization of such sleep schedules is challenging in itself, a larger concern is that arbitrarily long sleep periods can reduce the responsiveness and effectiveness of the sensors. In applications where it is critical that certain events in the environment be detected and reported rapidly, the latency induced by sleep schedules must be kept within strict bounds, even in the presence of network congestion.
3. **Robustness:** The vision of wireless sensor networks is to provide large-scale, yet fine-grained coverage. This motivates the use of large numbers of inexpensive devices. However, inexpensive devices can often be unreliable and prone to failures. Rates of device failure will also be high whenever the sensor devices are deployed in harsh or hostile environments. Protocol designs must therefore have built-in mechanisms to provide robustness. It is important to ensure that the global performance of the system is not sensitive to individual device failures. Further, it is often desirable that the performance of the system degrade as gracefully as possible with respect to component failures.
4. **Synergy:** Moore's law-type advances in technology have ensured that device capabilities in terms of processing power, memory, storage, radio transceiver performance, and even accuracy of sensing improve rapidly (given a fixed cost). However, if economic considerations dictate that the cost per node be reduced drastically from hundreds of dollars to less than a few cents, it is possible that the capabilities of individual nodes will remain constrained to some extent. The challenge is therefore to design synergistic protocols, which ensure that the system as a whole is more capable than the sum of the capabilities of its individual components. The protocols must provide an efficient collaborative use of storage, computation, and communication resources.
5. **Scalability:** For many envisioned applications, the combination of fine-granularity sensing and large coverage area implies that wireless sensor

networks have the potential to be extremely large scale (tens of thousands, perhaps even millions of nodes in the long term). Protocols will have to be inherently distributed, involving localized communication, and sensor networks must utilize hierarchical architectures in order to provide such scalability. However, visions of large numbers of nodes will remain unrealized in practice until some fundamental problems, such as failure handling and *in-situ* reprogramming, are addressed even in small settings involving tens to hundreds of nodes. There are also some fundamental limits on the throughput and capacity that impact the scalability of network performance.

6. **Heterogeneity:** There will be a heterogeneity of device capabilities (with respect to computation, communication, and sensing) in realistic settings. This heterogeneity can have a number of important design consequences. For instance, the presence of a small number of devices of higher computational capability along with a large number of low-capability devices can dictate a two-tier, cluster-based network architecture, and the presence of multiple sensing modalities requires pertinent sensor fusion techniques. A key challenge is often to determine the right combination of heterogeneous device capabilities for a given application.
7. **Self-configuration:** Because of their scale and the nature of their applications, wireless sensor networks are inherently *unattended* distributed systems. Autonomous operation of the network is therefore a key design challenge. From the very start, nodes in a wireless sensor network have to be able to configure their own network topology; localize, synchronize, and calibrate themselves; coordinate inter-node communication; and determine other important operating parameters.
8. **Self-optimization and adaptation:** Traditionally, most engineering systems are optimized *a priori* to operate efficiently in the face of expected or well-modeled operating conditions. In wireless sensor networks, there may often be significant uncertainty about operating conditions prior to deployment. Under such conditions, it is important that there be in-built mechanisms to autonomously learn from sensor and network measurements collected over time and to use this learning to continually improve performance. Also, besides being uncertain *a priori*, the environment in which the sensor network operates can change drastically over time. WSN protocols should also be able to adapt to such environmental dynamics in an online manner.
9. **Systematic design:** As we shall see, wireless sensor networks can often be highly application specific. There is a challenging tradeoff between (a) *ad hoc*, narrowly applicable approaches that exploit application-specific characteristics to offer performance gains and (b) more flexible, easy-to-generalize

design methodologies that sacrifice some performance. While performance optimization is very important, given the severe resource constraints in wireless sensor networks, systematic design methodologies, allowing for reuse, modularity, and run-time adaptation, are necessitated by practical considerations.

10. **Privacy and security:** The large scale, prevalence, and sensitivity of the information collected by wireless sensor networks (as well as their potential deployment in hostile locations) give rise to the final key challenge of ensuring both privacy and security.

1.5 Organization

This book is organized in a bottom–up manner. Chapter 2 addresses tools, techniques, and metrics pertinent to network deployment. Chapter 3 and Chapter 4 present techniques for spatial localization and temporal synchronization respectively. Chapter 5 addresses a number of issues pertaining to wireless characteristics, including models for link quality, interference, and radio energy. Algorithms for medium-access and radio sleep scheduling for energy conservation are described in Chapter 6. Topology control techniques based on sleep–active transitions are described in Chapter 7. Mechanisms for energy-efficient and robust routing are discussed in Chapter 8, while Chapter 9 presents concepts and techniques for data-centric routing and querying in wireless sensor networks. Chapter 10 covers issues pertinent to congestion control and transport-layer quality of service. Finally, we present concluding comments in Chapter 11, along with a brief survey of some important further topics.

2

Network deployment

2.1 Overview

The problem of *deployment* of a wireless sensor network could be formulated generically as follows: given a particular application context, an operational region, and a set of wireless sensor devices, how and where should these nodes be placed?

The network must be deployed keeping in mind two main objectives: coverage and connectivity. *Coverage* pertains to the application-specific quality of information obtained from the environment by the networked sensor devices. *Connectivity* pertains to the network topology over which information routing can take place. Other issues, such as equipment costs, energy limitations, and the need for robustness, should also be taken into account.

A number of basic questions must be considered when deploying a wireless sensor network:

1. **Structured versus randomized deployment:** Does the network involve (a) structured placement, either by hand or via autonomous robotic nodes, or (b) randomly scattered deployment?
2. **Over-deployment versus incremental deployment:** For robustness against node failures and energy depletion, should the network be deployed *a priori* with redundant nodes, or can nodes be added or replaced incrementally when the need arises? In the former case, sleep scheduling is desirable to extend network lifetime, a topic we will treat in Chapter 7.
3. **Network topology:** Is the network topology going to be a simple star topology, or a grid, or an arbitrary multi-hop mesh, or a two-level cluster hierarchy? What kind of robust connectivity guarantees are desired?