# 1. Hardware components

## 1.1 Sensor node hardware overview

When choosing the hardware components for a wireless sensor node, evidently the application's requirements play a decisive factor with regard mostly to size, costs, and energy consumption of the nodes – communication and computation facilities as such are often considered to be of acceptable quality, but the trade-offs between features and costs is crucial. In some extreme cases, an entire sensor node should be smaller than 1 cc, weigh (considerably) less than 100 g, be substantially cheaper than US$1, and dissipate less than 100 $\mu$W. In even more extreme visions, the nodes are sometimes claimed to have to be reduced to the size of grains of dust. In more realistic applications, the mere size of a node is not so important; rather, convenience, simple power supply, and cost are more important .

These diversities notwithstanding, a certain common trend is observable in the literature when looking at typical hardware platforms for wireless sensor nodes. While there is certainly not a single standard available, nor would such a standard necessarily be able to support all application types, this section will survey these typical sensor node architectures. In addition, there are a number of research projects that focus on shrinking any of the components in size, energy consumption, or costs, based on the fact that custom off-the-shelf components do currently not live up to some of the more stringent application requirements. But as this book focuses on the networking aspects of WSNs, these efforts are not discussed here.

A basic sensor node comprises five main components (Figure 2.1):

**Controller** A controller to process all the relevant data, capable of executing arbitrary code.

**Memory** Some memory to store programs and intermediate data; usually, different types of memory are used for programs and data.

**Sensors and actuators** The actual interface to the physical world: devices that can observe or control physical parameters of the environment.

**Communication** Turning nodes into a network requires a device for sending and receiving infor-mation over a wireless channel.

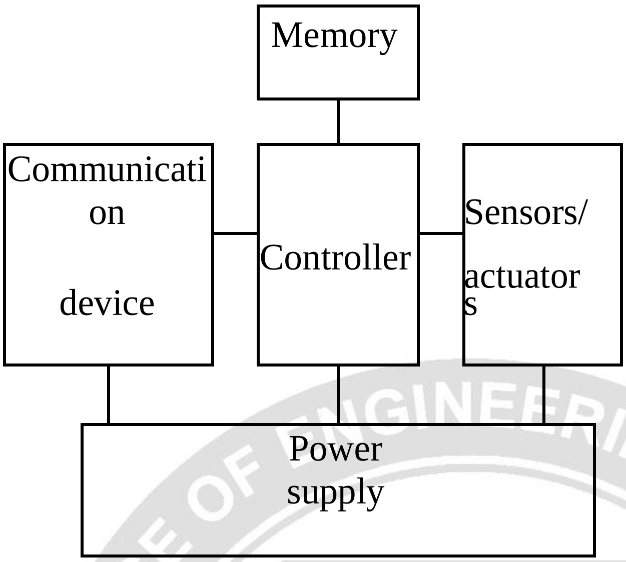


**Figure 2.1** Overview of main sensor node hardware components

---

**Power supply** As usually no tethered power supply is available, some form of batteries are neces-sary to provide energy. Sometimes, some form of recharging by obtaining energy from the environment is available as well (e.g. solar cells).

Each of these components has to operate balancing the trade-off between as small an energy consumption as possible on the one hand and the need to fulfill their tasks on the other hand. For example, both the communication device and the controller should be turned off as long as possible. To wake up again, the controller could, for example, use a preprogrammed timer to be reactivated after some time. Alternatively, the sensors could be programmed to raise an interrupt if a given event occurs – say, a temperature value exceeds a given threshold or the communication device detects an incoming transmission.
Supporting such alert functions requires appropriate interconnection between individual compo-nents. Moreover, both control and data information have to be exchanged along these interconnec-tions. This interconnection can be very simple – for example, a sensor could simply report an analog value to the controller – or it could be endowed with some intelligence of its own, preprocessing sensor data and only waking up the main controller if an actual event has been detected – for example, detecting a threshold crossing for a simple temperature sensor. Such preprocessing can be highly customized to the specific sensor yet remain simple enough to run continuously, resulting in improved energy efficiency [26].

## 1.2 Controller

### Microcontrollers versus microprocessors, FPGAs, and ASICs

The controller is the core of a wireless sensor node. It collects data from the sensors, processes this data, decides when and where to send it, receives data from other sensor nodes, and decides on the actuator's behavior. It has to execute various programs, ranging from time-critical signal processing and communication protocols to application programs; it is the Central Processing Unit (CPU) of the node.

Such a variety of processing tasks can be performed on various controller architectures, representing trade-offs between flexibility, performance, energy efficiency, and costs.

One solution is to use general-purpose processors, like those known from desktop computers. These processors are highly overpowered, and their energy consumption is excessive. But simpler processors do exist, specifically geared toward usage in embedded systems. These processors are commonly referred as **microcontrollers**. Some of the key characteristics why these microcontrollers are particularly suited to embedded systems are their flexibility in connecting with other devices (like sensors), their instruction set amenable to time-critical signal processing, and their typically low power consumption; they are also convenient in that they often have memory built in. In addition, they are freely programmable and hence very flexible. Microcontrollers are also suitable for WSNs since they commonly have the possibility to reduce their power consumption by going into **sleep states** where only parts of the controller are active; details vary considerably between different controllers. Details regarding power consumption and energy efficiency are discussed in Section 2.2. One of the main differences to general-purpose systems is that microcontroller-based systems usually do not feature a memory management unit, somewhat limiting the functionality of memory – for example, protected or virtual memory is difficult, if not impossible, to achieve.

Single-node architecture

. Hence, these advantages of a DSP are typically not required in a WSN node and they are usually not used. Another option for the controller is to depart from the high flexibility offered by a (fairly general-purpose) microcontroller and to use Field-Programmable Gate Arrays (FPGAs) or Application-Specific Integrated Circuits (ASICs) instead. An FPGA can be reprogrammed (or rather recon-figured) "in the field" to adapt to a changing set of requirements; however, this can take time and energy – it is not practical to reprogram an FPGA at the same frequency as a microcontroller could change between different programs. An ASIC is a specialized processor, custom designed for a given application such as, for example, high-speed routers and switches. The typical trade-off here is loss of flexibility in return for a considerably better energy efficiency and performance. On the other hand, where a microcontroller requires software development, ASICs provide the same functionality in hardware, resulting in potentially more costly hardware development.

For a dedicated WSN application, where the duties of a the sensor nodes do not change over lifetime and where the number of nodes is big enough to warrant the investment in ASIC development, they can be a superior solution. At the current stage of WSN technology, however, the bigger flexibility and simpler usage of microcontrollers makes them the generally preferred solution. However, this is not necessarily the final solution as "convenient programmability over several orders of energy consumption and data processing requirements is a worthy research goal" . In addition, splitting processing tasks between some low-level, fixed functionality put into a very energy-efficient ASIC and high-level, flexible, relatively rarely invoked processing on a micro controller is an attractive design and research option.

**For the remainder of this book, a microcontroller-based architecture is assumed.**

**Some examples for microcontrollers**

Microcontrollers that are used in several wireless sensor node prototypes include the Atmel processor or Texas Instrument's MSP 430. In older prototypes, the Intel Strong Arm processors have also been used, but this is no longer considered as a practical option; it is included here for the sake of completeness. Nonetheless, as the principal properties of these processors and controllers are quite similar, conclusions from these earlier research results still hold to a large degree.

**Intel Strong ARM**

The Intel Strong ARM [379] is, in WSN terms, a fairly high-end processor as it is mostly geared toward handheld devices like PDAs. The SA-1100 model has a 32-bit Reduced Instruction Set Computer (RISC) core, running at up to 206 MHz.

**Texas Instruments MSP 430**

Texas Instrument provides an entire family of microcontrollers under the family designation MSP 430 Unlike the Strong ARM, it is explicitly intended for embedded applications. Accordingly, it runs a 16-bit RISC core at considerably lower clock frequencies (up to 4 MHz) but comes with a wide range of interconnection possibilities and an instruction set amenable to easy handling of peripherals of different kinds. It features a varying amount of on-chip RAM (sizes are 2 – 10 kB), several 12-bit analog/digital converters, and a real-time clock. It is certainly powerful enough to handle the typical computational tasks of a typical wireless sensor node (possibly with the exception of driving the radio front end, depending on how it is connected – bit or byte interface – to the controller).

---

**Atmel ATmega**

The Atmel ATmega 128L [28] is an 8-bit microcontroller, also intended for usage in embedded applications and equipped with relevant external interfaces for common peripherals.

### 2.1.3 Memory

The memory component is fairly straightforward. Evidently, there is a need for Random Access Memory (RAM) to store intermediate sensor readings, packets from other nodes, and so on. While RAM is fast, its main disadvantage is that it loses its content if power supply is interrupted. Program code can be stored in Read-Only Memory (ROM) or,

more typically, in Electrically Erasable Programmable Read-Only Memory (EEPROM) or flash memory (the later being similar to EEPROM but allowing data to be erased or written in blocks instead of only a byte at a time). Flash memory can also serve as intermediate storage of data in case RAM is insufficient or when the power supply of RAM should be shut down for some time. The long read and write access delays of flash memory should be taken into account, as well as the high required energy.
Correctly dimensioning memory sizes, especially RAM, can be crucial with respect to manufacturing costs and power consumption. However, even general rules of thumbs are difficult to give as the memory requirements are very much application dependent.

### 2.1.4 Communication device

### Choice of transmission medium

The communication device is used to exchange data between individual nodes. In some cases, wired communication can actually be the method of choice and is frequently applied in many sensor network like settings (using field buses like Profibus, LON, CAN, or others). The communication devices for these networks are custom off-the-shelf components.
The case of wireless communication is considerably more interesting. The first choice to make is that of the transmission medium – the usual choices include radio frequencies, optical communication, and ultrasound; other media like magnetic inductance are only used in very specific cases. Of these choices, Radio Frequency (RF)-based communication is by far the most relevant one as it best fits the requirements of most WSN applications: It provides relatively long range and high data rates, acceptable error rates at reasonable energy expenditure, and does not require line of sight between sender and receiver. Thus, RF-based communication and transceiver will receive the lion share of attention here; other media are only treated briefly at the end of this section.
For a practical wireless, RF-based system, the carrier frequency has to be carefully chosen. Chapter 4 contains a detailed discussion; for the moment, suffice it to say that wireless sensor networks typically use communication frequencies between about 433 MHz and 2.4 GHz.

### Transceivers

For actual communication, both a transmitter and a receiver are required in a sensor node. The essential task is to convert a bit stream coming from a microcontroller (or a sequence of bytes or frames) and convert them to and from radio waves. For practical purposes, it is usually convenient to use a device that combines these two tasks in a single entity. Such combined devices are called **transceivers**. Usually, half-duplex operation is realized since transmitting and receiving at the same time on a wireless medium is impractical in most cases (the receiver would only hear the own transmitter anyway).

A range of low-cost transceivers is commercially available that incorporate all the circuitry required for transmitting and receiving – modulation, demodulation, amplifiers, filters, mixers, and so on. For a judicious choice, the transceiver's tasks and its main characteristics have to be understood.

**Transceiver tasks and characteristics**

To select appropriate transceivers, a number of characteristics should be taken into account. The most important ones are:

**Service to upper layer** A receiver has to offer certain services to the upper layers, most notably to the Medium Access Control (MAC) layer. Sometimes, this service is **packet oriented**; sometimes, a transceiver only provides a **byte interface** or even only a **bit interface** to the microcontroller.

In any case, the transceiver must provide an interface that somehow allows the MAC layer to initiate frame transmissions and to hand over the packet from, say, the main memory of the sensor node into the transceiver (or a byte or a bit stream, with additional processing required on the microcontroller). In the other direction, incoming packets must be streamed into buffers accessible by the MAC protocol.

**Power consumption and energy efficiency** The simplest interpretation of energy efficiency is the energy required to transmit and receive a single bit. Also, to be suitable for use in WSNs, transceivers should be switchable between different states, for example, active and sleeping. The idle power consumption in each of these states and during switching between them is very important – details are discussed in Section 2.2.

**Carrier frequency and multiple channels** Transceivers are available for different carrier frequencies; evidently, it must match application requirements and regulatory restrictions. It is often useful if the transceiver provides several carrier frequencies ("channels") to choose from, helping to alleviate some congestion problems in dense networks. Such channels or "sub-bands" are relevant, for example, for certain MAC protocols (FDMA or multichannel CSMA/ ALOHA techniques, see Chapter 5).

**State change times and energy** A transceiver can operate in different modes: sending or receiv-ing, use different channels, or be in different power-safe states. In any case, the time and the energy required to change between two such states are important figures of merit. The turnaround time between sending and receiving, for example, is important for various medium access protocols (see Chapter 5).

**Data rates** Carrier frequency and used bandwidth together with modulation and coding determine the gross data rate. Typical values are a few tens of kilobits per second – considerably less than in broadband wireless communication, but usually sufficient for WSNs. Different data rates can be achieved, for example, by using different modulations or changing the symbol rate.

**Modulations** The transceivers typically support one or several of on/off-keying, ASK, FSK, or similar modulations. If several modulations are available, it is convenient for experiments if they are selectable at runtime even though, for real deployment, dynamic switching between modulations is not one of the most discussed options.

**Coding** Some transceivers allow various coding schemes to be selected.

**Transmission power control** Some transceivers can directly provide control over the transmission power to be used; some require some external circuitry for that purpose. Usually, only a

discrete number of power levels are available from which the actual transmission power can be chosen. Maximum output power is usually determined by regulations.

**Noise figure** The **noise figure** NF of an element is defined as the ratio of the Signal-to-Noise Ratio (SNR) ratio $SNR_I$ at the input of the element to the SNR ratio $SNR_O$ at the element's output:

$$NF = \frac{SNR_I}{SNR_O}$$

It describes the degradation of SNR due to the element's operation and is typically given in dB:

$$NF\ dB = SNR_I\ dB - SNR_O\ dB$$

**Gain** The **gain** is the ratio of the output signal power to the input signal power and is typically given in dB. Amplifiers with high gain are desirable to achieve good energy efficiency.

**Power efficiency** The **efficiency** of the radio front end is given as the ratio of the radiated power to the overall power consumed by the front end; for a power amplifier, the efficiency describes the ratio of the output signal's power to the power consumed by the overall power amplifier.

**Receiver sensitivity** The **receiver sensitivity** (given in dBm) specifies the minimum signal power at the receiver needed to achieve a prescribed $E_b/N_0$ or a prescribed bit/packet error rate. Better sensitivity levels extend the possible range of a system.

**Range** While intuitively the range of a transmitter is clear, a formal definition requires some care. The range is considered in absence of interference; it evidently depends on the maximum transmission power, on the antenna characteristics, on the attenuation caused by the environ-ment, which in turn depends on the used carrier frequency, on the modulation/coding scheme that is used, and on the bit error rate that one is willing to accept at the receiver. It also depends on the quality of the receiver, essentially captured by its sensitivity. Typical values are difficult to give here, but prototypes or products with ranges between a few meters and several hundreds of meters are available.

**Blocking performance** The blocking performance of a receiver is its achieved bit error rate in the presence of an interferer. More precisely, at what power level can an interferer (at a fixed distance) send at a given offset from the carrier frequency such that target BER can still be met? An interferer at higher frequency offsets can be tolerated at large power levels. Evidently, blocking performance can be improved by interposing a filter between antenna and transceiver.

An important special case is an adjacent channel interferer that transmits on neighboring frequencies. The adjacent channel suppression describes a transceiver’s capability to filter out signals from adjacent frequency bands (and thus to reduce adjacent channel interference) has a direct impact on the observed Signal to Interference and Noise Ratio (SINR).

**Out of band emission** The inverse to adjacent channel suppression is the out of band emission of a transmitter. To limit disturbance of other systems, or of the WSN itself in a multichannel setup, the transmitter should produce as little as possible of transmission power outside of its prescribed bandwidth, centered around the carrier frequency.

**Carrier sense and RSSI** In many medium access control protocols, sensing whether the wireless channel, the carrier, is busy (another node is transmitting) is a critical information. The

receiver has to be able to provide that information. The precise semantics of this carrier-sense signal depends on the implementation. For example, the IEEE 802.15.4 standard [468] distinguishes the following modes:

The received energy is above threshold; however, the underlying signal does not need to comply with the modulation and spectral characteristics.

A carrier has been detected, that is, some signal which complies with the modulation.

Carrier detected and energy is present.

Also, the signal strength at which an incoming data packet has been received can provide useful information (e.g. a rough estimate about the distance from the transmitter assuming the transmission power is known); a receiver has to provide this information in the Received Signal Strength Indicator (RSSI).

**Frequency stability** The **frequency stability** denotes the degree of variation from nominal center frequencies when environmental conditions of oscillators like temperature or pressure change. In extreme cases, poor frequency stability can break down communication links, for example, when one node is placed in sunlight whereas its neighbor is currently in the shade.

**Voltage range** Transceivers should operate reliably over a range of supply voltages. Otherwise, inefficient voltage stabilization circuitry is required.

Transceivers appropriate for WSNs are available from many manufacturers. Usually, there is an entire family of devices to choose from, for example, customized to different regulatory restrictions on carrier frequency in Europe and North America. Currently popular product series include the RFM TR 1001, the Chipcon CC 1000 and CC 2420 (as one of the first IEEE 802.15.4 compliant models), and the Infineon

For simple transceivers, the additional cost of providing such an identifier is relatively high with respect to the device's total costs, and thus, unique identifiers cannot be
- relied upon to be present in all devices. The availability of such device identifiers is very useful in many communication protocols and their absence will have considerable consequences for protocol design.

**Transceiver structure**

A fairly common structure of transceivers is into the Radio Frequency (RF) front end and the baseband part:

the **radio frequency front end** performs analog signal processing in the actual radio frequency band, whereas

the **baseband processor** performs all signal processing in the digital domain and communicates with a sensor node's processor or other digital circuitry.

Between these two parts, a frequency conversion takes place, either directly or via one or several Intermediate Frequencys (IFs). The boundary between the analog and the digital domain is constituted by Digital/Analog Converters (DACs) and Analog/Digital Converters (ADCs).

The **RF front end** performs analog signal processing in the actual radio frequency band, for example in the 2.4 GHz Industrial, Scientific, and Medical (ISM) band; Some important elements of an RF front ends architecture are sketched in Figure 2.2:

The Power Amplifier (PA) accepts up converted signals from the IF or baseband part and amplifies them for transmission over the antenna.

The range of powers of the incoming signals varies from very weak signals from nodes close to the reception boundary to strong signals from nearby nodes; this range can be up to 100 dB. Without management actions, the LNA is active all the time and can consume a significant fraction of the transceiver's energy.

Elements like local oscillators or voltage-controlled oscillators and mixers are used for frequency conversion from the RF spectrum to intermediate frequencies or to the baseband. The incoming signal at RF frequencies $f_{RF}$ is multiplied in a mixer with a fixed-frequency signal from the local oscillator (frequency $f_{LO}$). The resulting intermediate-frequency signal has frequency $f_{LO} - f_{RF}$. Depending on the RF front end architecture, other elements like filters are also present.

**Transceiver operational states**

Many transceivers can distinguish four operational states [670]:

Radio frontend

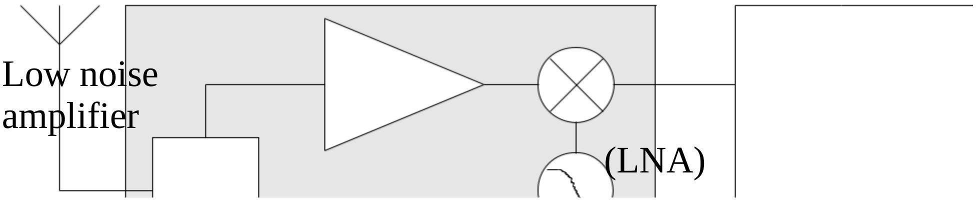


Antenna
interface

Power
amplifier
(PA)

Frequency Conversion
Intermediate frequency and baseband processing

**Figure 2.2** RF front end

**Transmit** In the **transmit state**, the transmit part of the transceiver is active and the antenna radiates energy.

**Receive** In the **receive state** the receive part is active.

**Idle** A transceiver that is ready to receive but is not currently receiving anything is said to be in an **idle state**. In this idle state, many parts of the receive circuitry are active, and others can be switched off. For example, in the synchronization circuitry, some elements concerned with acquisition are active, while those concerned with tracking can be switched off and activated only when the acquisition has found something. MYERS et al. [580] also discuss techniques for switching off parts of the acquisition circuitry for IEEE 802.11 transceivers. A major source of power dissipation is **leakage**.

**Sleep** In the **sleep state**, significant parts of the transceiver are switched off. There are transceivers offering several different sleep states, see reference [580] for a discussion of sleep states for IEEE 802.11 transceivers. These sleep states differ in the amount of circuitry switched off and in the associated **recovery times** and **startup energy** [855]. For example, in a complete power down of the transceiver, the startup costs include a complete initialization as well as configuration of the radio, whereas in "lighter" sleep modes, the clock driving certain transceiver parts is throttled down while configuration and operational state is remembered.

The sensor node's protocol stack and operating software must decide into which state the trans-ceiver is switched, according to the current and anticipated communications needs. One problem complicating this decision is that the operation of state changes also dissipate power [670]. For example, a transceiver waking up from the sleep mode to the transmit mode requires some startup time and startup energy, for example, to ramp up phase-locked loops or voltage-controlled oscilla-tors. During this startup time, no transmission or reception of data is possible [762]. The problem of scheduling the node states (equivalently: switching on and off node/transceiver components) so as to minimize average power consumption (also called **power management**) is rather complex, an in-depth treatment can be found in reference [85], and a further reference is [741].

### Advanced radio concepts

Apart from these basic transceiver concepts, a number of advanced concepts for radio communi-cation are the objectives of current research. Three of them are briefly summarized here.

### Wakeup radio

Looking at the transceiver concepts described above, one of the most power-intensive operations is waiting for a transmission to come in, ready to receive it. During this time, the receiver circuit must be powered up so that the wireless channel can be observed, spending energy without any immediate benefit.

While it seems unavoidable to provide a receiver with power during the actual reception of a packet, it would be desirable not to have to invest power while the node is only waiting for a packet to come in. A receiver structure is necessary that does not need power but can detect when a packet starts to arrive. To keep this specialized receiver simple, it suffices for it to raise an event to notify other components of an incoming packet; upon such an event, the main receiver can be turned on and perform the actual reception of the packet.

Such receiver concepts are called **wakeup receivers** : Their only purpose is to wake up the main receiver without needing (a significant amount of) power to do so – ZHONG et al. state a target power consumption of less than 1 □W. In the simplest case, this wakeup would happen for every packet; a more sophisticated version would be able to decide,

- using proper address information at the start of the packet, whether the incoming packet is actually destined for this node and only then wake up the main receiver.

Such wakeup receivers are tremendously attractive as they would do away with one of the main problems of WSNs: the need to be permanently able to receive in a network with low average traffic. It would considerably simplify a lot of the design problems of WSNs, in particular of the medium access control – Section 5.2.4 will discuss these aspects and some ensuing problems in more detail. Unfortunately, so far the realization of a reliable, well-performing wakeup receiver has not been achieved yet.

**Spread-spectrum transceivers**

Simple transceiver concepts, based on modulations like Amplitude Shift Keying (ASK) or Fre-quency Shift Keying (FSK), can suffer from limited performance, especially in scenarios with a lot of interference. To overcome this limitation, the use of spread-spectrum transceivers has been proposed by some researchers [155, 281] . These transceivers, however, suffer mostly from complex hardware and consequently higher prices, which has prevented them from becoming a mainstream concept for WSNs so far. Section 4.2.5 presents details.

**Ultra wideband communication**

Ultra Wide Band (UWB) communication is a fairly radical change from conventional wireless communication as outlined above. Instead of modulating a digital signal onto a carrier frequency, a very large bandwidth is used to directly transmit the digital sequence as very short impulses (to form nearly rectangular impulses requires considerable bandwidth, because of which this con-cept is not used traditionally)  Accordingly, these impulses occupy a large spectrum starting from a few Hertz up to the range of several GHz. The challenge is to syn-chronize sender and receiver sufficiently (to an accuracy of trillionth of seconds) so that the impulses can be correctly detected. A side effect of precisely timed impulses is that UWB is fairly resistant to multipath fading which can be a serious obstacle for carrier-based radio communication.

As one concrete example, consider a time-hopping Pulse Position Modulation (PPM) proposed as combined modulation and multiple access scheme by WIN and SCHOLTZ [. For each symbol, a number of pulses are transmitted with almost periodic spacing. The deviations from the periodicity encode both the modulation as well as the transmitting user.

For a communication system, the effect is that a very high data rate can be realized over short distances; what is more, UWB communication can relatively easily penetrate obstacles such as doors, which are impermeable to narrowband radio waves. For a WSN, the high data rate is not strictly necessary but can be leveraged to reduce the on-time of the transceivers. The nature of UWB also allows to precisely measure distances (with claimed precision of centimeters).

These desirable features of UWB communication have to be balanced against the difficulties of building such transceivers at low-cost and low-power consumption. More precisely, an UWB transmitter is actually relatively simple since it does not need oscillators or related circuitry found in transmitters for a carrier-frequency-based transmitter. The receivers, on the other hand, require complex timing synchronization. As of this writing, UWB transceivers have not yet been used in prototypes for wireless sensor nodes.

One of the best sources of information about UWB in WSN might be the documents of the IEEE 802.15.4a study group, which looks at UWB as an alternative physical layer for the IEEE 802.15.4 standard for short-range, low bitrate wireless communication. Some references to start from are [82, 187, 566, 603, 884]. A comparison between UWB and Direct Sequence Spread Spectrum (DSSS) technologies for sensor networks has been made in [939], under the assumption of an equal bandwidth for both types of systems.

**Nonradio frequency wireless communication**

While most of the wireless sensor network work has focused on the use of radio waves as communication media, other options exists. In particular, optical communication and ultrasound communication have been considered as alternatives.

**Optical**

the use of optical links between sensor nodes. Its main advantage is the very small energy per bit required for both generating and detecting optical light – simple Light-Emitting Diodes (LEDs) are good examples for high-efficiency senders. The required circuitry for an optical transceiver is also simpler and the device as a whole can be smaller than the radio frequency counterpart. Also, communication can take place concurrently with only negligible interference. The evident disadvantage, however, is that communicating peers need to have a line of sight connection and that optical communication is more strongly influenced by weather conditions.

As a case in point, consider the so-called “corner-cube reflector”: three mirrors placed at right angles to each other in a way that each beam of light directed at it is reflected back to its source (as long as it comes from a cone centered around the main diagonal of the cube) – an example for such a structure is shown in Figure 2.3. This reflection property holds only as long as the mirrors are exactly at right angles. When one the mirrors is slightly moved, a signal can be modulated onto an incoming ray of light, effectively transmitting information back to the sender. In fact, data rates up to 1 kb/s have been demonstrated using such a device. Its main advantage is that the mechanical movement of one such mirror only takes very little energy, compared to actually generating a beam of light or even a radio wave. Hence, a passive readout of sensor nodes can be done very energy

**Figure 2.3** Example of a corner-cube reflector for optical communication [168]. Reproduced by permission of IEEE

efficiently over long distances as long as the reader has enough power to produce the laser beam (up to 150 m have been demonstrated using a 5 mW laser).

**Ultrasound**

Both radio frequency and optical communication are suitable for open-air environments. In some application scenarios, however, sensor nodes are used in environments where radio or optical communication is not applicable because these waves do not penetrate the surrounding medium. One such medium is water, and an application scenario is the surveillance of marine ground floor erosion to help in the construction of offshore wind farms. Sensors are deployed on the marine ground floor and have to communicate amongst themselves. In such an underwater environment, ultrasound is an attractive communication medium as it travels relatively long distances at comparably low power.

**Some examples of radio transceivers**

To complete this discussion of possible communication devices, a few examples of standard radio transceivers that are commonly used in various WSN prototype nodes should be briefly described. All these transceivers are in fact commodity, off-the-shelf items available via usual distributors. They are all single-chip solutions, integrating transmitter and receiver functionality, requiring only a small number of external parts and have a fairly low-power consumption. In principle, similar equipment is available from a number of manufacturers – as can be expected, there is not one "best product" available, but each of them has particular advantages and disadvantages.

**RFM TR1000 family**

The TR1000 family of radio transceivers from RF Monolithics[2] is available for the 916 MHz and 868 MHz frequency range. It works in a 400 kHz wide band centered at, for example, 916.50 MHz. It is intended for short-range radio communication with up to 115.2 kbps. The modulation is either on-off-keying (at a maximum rate of 30 kbps) or ASK; it also provides a dynamically tunable output power. The maximum radiated power is given in the data sheet [690] as 1.5 dBm, ≈ 1.4 mW, whereas in the Mica motes a number of 0.75 mW is given . The transceiver offers received signal strength information. It is attractive because of its low-power consumption in both send and receive modes and especially in sleep mode. Details about parameters and configurations can be found in the data sheet.

**Hardware accelerators (Mica motes)**

The Mica motes use the RFM TR1000 transceiver and contain also a set of **hardware accelerators**. On the one hand, the transceiver offers a very low-level interface, giving the microcontroller tight control over frame formats, MAC protocols, and so forth. On the other hand, framing and MAC can be very computation intensive, for example, for computing checksums, for making bytes out of serially received bits or for detecting Start Frame Delimiters (SFDs) in a stream of symbols. The hardware accelerators offer some of these primitive computations in hardware, right at the disposal of the microcontroller.

**Chipcon CC1000 and CC2420 family**

Chipcon offers a wide range of transceivers that are appealing for use in WSN hardware. To name but two examples: The CC1000 operates in a wider frequency range, between 300 and 1000 MHz,

programmable in steps of 250 Hz. It uses FSK as modulation, provides RSSI, and has programmable output power. An interesting feature is the possibility to compensate for crystal temperature drift. It should also be possible to use it in frequency hopping protocols. Details can be found in the data sheet.

The CC2420 is a more complicated device. It implements the physical layer as prescribed by the IEEE 802.15.4 standard with the required support for this standard's MAC protocol. In fact, the company claims that this is the first commercially available single-chip transceiver for IEEE 802.15.4. As a consequence of implementing this standard, the transceiver operates in the 2.4 GHz band and features the required DSSS modem, resulting in a data rate of 250 kbps. It achieves this at still relatively low-power consumption, although not quite on par with the simpler transceivers described so far.

**Infineon TDA 525x family**

The Infineon TDA 525x family provides flexible, single-chip, energy-efficient transceivers. The TDA , as an example, is a 868 – 870 MHz transceiver providing both ASK and FSK modulation, it has a highly efficient power amplifier, RSSI information, a tunable crystal oscillator, an onboard data filter, and an intelligent power-down feature. One of the interesting features is a self-polling mechanism, which can very quickly determine data rate. Compared to some other transceiver, it also has an excellent blocking performance that makes it quite resistant to interference.

**IEEE 802.15.4/Ember EM2420 RF transceiver**

The IEEE 802.15.4 low-rate Wireless Personal Area Network (WPAN) works in three differ-ent frequency bands and employs a DSSS scheme. Some basic data can be found in Table 2.1. For one particular RF front-end design, the Ember[4] EM2420 RF Transceiver , some numbers on power dissipation are available. For a radiated power of −0.5 dBm (corresponding to ≈0.9 mW) and with a supply voltage of 3.3 V, the transmit mode draws a current of 22.7 mA, correspond-ing to ≈74.9 mW, whereas in the receive mode, 25.2 mA current are drawn, corresponding to ≈83.2 mW. In the sleep mode, only 12 □A are drawn.

In all bands, DSSS is used. In the 868 MHz band, only a single channel with a data rate of 20 kbps is available, in the 915 MHz band ten channels of 40kbps each and in the 2.4 GHz band 16 channels of 250 kbps are available. In the lower two bands, the chips are Binary Phase Shift Keying (BPSK)-modulated, and the data symbols are encoded differentially. A pseudo noise sequence of 15 chips is used for every bit. The modulation scheme in the 2.4 GHz band is a little

**Table 2.1** The different PHY’s of the IEEE 802.15.4 standard [468]. Reproduced by permission of IEEE

| Band | 868 MHz | 915 MHz | 2.4 GHz |
|---|---|---|---|
| Frequency [MHz] | 868–868.6 | 902–928 | 2400–2483.5 |
| Chip rate [kchips/s] | 300 | 600 | 2000 |
| # of channels | 1 | 10 | 16 |
| Modulation | BPSK | BPSK | O-QPSK |

| | | | |
|---|---|---|---|
| Data rate [kb/s] | 20 | 40 | 250 |
| Symbol rate [ksymbols/s] | 20 | 40 | 62.5 |
| Symbol type | binary | binary | 16-ary orthogonal |

bit more complicated. As can be observed from the table, a channel symbol consists of four user bits. These 16 different symbol values are distinguished by using 16 different nearly orthogonal pseudorandom chip sequences. The resulting chip sequence is then modulated using a modulation scheme called *offset* -Quaternary Phase Shift Keying (QPSK). Some of the design rationale for this modulation scheme

**National Semiconductor LMX3162**

The radio hardware of the $\mu$AMPS-1 node] consists of a digital baseband processor implemented on an FPGA, whereas for the RF front end, a (now obsolete) National Semiconductor LMX3162 transceiver [588] is used. The LMX3162 operates in the 2.4 GHz band and offers six different radiated power levels from 0 dBm up to 20 dBm. To transmit data, the baseband processor can control an externally controllable Voltage-Controlled Oscillator (VCO). The main components of the RF front end (phase-lock loop, transmit and receive circuitry) can be shut off. The baseband processor controls the VCO and also provides timing information to a TDMA-based MAC protocol (see Chapter 5). For data transmission, FSK with a data rate of 1 Mbps is used.

**Conexant RDSSS9M**

The WINS sensor node of Rockwell[5] carries a Conexant RDSSS9M transceiver, consisting of the RF part working in the ISM band between 902 and 928 MHz and a microcontroller (a 65C02) responsible for processing DSSS signals with a spreading factor of 12 bits per chip. The data rate is 100 kbps. The RF front end offers radiated power levels of 1 mW, 10 mW and 100 mW. A number of 40 sub-bands are available, which can be freely selected. The microcontroller implements portions of a MAC protocol also.

### 2.1.5 Sensors and actuators

Without the actual sensors and actuators, a wireless sensor network would be beside the point entirely. But as the discussion of possible application areas has already indicated, the possible range of sensors is vast. It is only possible to give a rough idea on which sensors and actuators can be used in a WSN.

**Sensors**

Sensors can be roughly categorized into three categories **Passive, omnidirectional sensors** These sensors can measure a physical quantity at the point of the sensor node without actually manipulating the environment by active probing – in this sense, they are passive. Moreover, some of these sensors actually are self-powered in the sense that they obtain the energy they need from the environment – energy is only needed to amplify their analog signal. There is no notion of "direction" involved in these measurements. Typical examples for such sensors include thermometer, light sensors, vibration, microphones, humidity, mechanical stress or tension in materials, chemical sensors sensitive for given substances, smoke detectors, air pressure, and so on.

**Passive, narrow-beam sensors** These sensors are passive as well, but have a well-defined notion of direction of measurement. A typical example is a camera, which can "take measurements" in a given direction, but has to be rotated if need be.

**Active sensors** This last group of sensors actively probes the environment, for example, a sonar or radar sensor or some types of seismic sensors, which generate shock waves by small

explosions. These are quite specific – triggering an explosion is certainly not a lightly under-taken action – and require quite special attention.

In practice, sensors from all of these types are available in many different forms with many individual peculiarities. Obvious trade-offs include accuracy, dependability, energy consumption, cost, size, and so on – all this would make a detailed discussion of individual sensors quite ineffective.

Overall, most of the theoretical work on WSNs considers passive, omnidirectional sensors. Narrow-beam-type sensors like cameras are used in some practical testbeds, but there is no real systematic investigation on how to control and schedule the movement of such sensors. Active sensors are not treated in the literature to any noticeable extent.

An assumption occasionally made in the literature is that each sensor node has a certain **area of coverage** for which it can reliably and accurately report the particular quantity that it is observing. More elaborately, a sensor detection model is used, relating the distance between a sensor and the to-be-detected event or object to a detection probability; an example for such a detection model is contained in references .

Strictly speaking, this assumption of a coverage area is difficult to justify in its simplest form. Nonetheless, it can be practically useful: It is often possible to postulate, on the basis of application-specific knowledge, some properties of the physical quantity under consideration, in particular, how quickly it can change with respect to distance. For example, temperature or air pressure are unlikely to vary very strongly within a few meters. Hence, allowing for some inevitable inaccuracies in the measurement, the maximum rate of changeover distance can be used to derive such a "coverage radius" within which the values of a single sensor node are considered "good enough". The precise mathematical tools for such a derivation are spatial versions of the sampling theorems.

**Actuators**

Actuators are just about as diverse as sensors, yet for the purposes of designing a WSN, they are a bit simpler to take account of: In principle, all that a sensor node can do is to open or close a switch or a relay or to set a value in some way. Whether this controls a motor, a light bulb, or some other physical object is not really of concern to the way communication protocols are designed. Hence, in this book, we shall treat actuators fairly summarily without distinguishing between different types.

In a real network, however, care has to be taken to properly account for the idiosyncrasies of different actuators. Also, it is good design practice in most embedded system applications to pair any actuator with a controlling sensor – following the principle to "never trust an actuator" .

### 2.1.6 Power supply of sensor nodes

For untethered wireless sensor nodes, the power supply is a crucial system component. There are essentially two aspects: First, storing energy and providing power in the required form; second, attempting to replenish consumed energy by "scavenging" it from some node-external power source over time.

Storing power is conventionally done using batteries. As a rough orientation, a normal AA battery stores about 2.2 – 2.5 Ah at 1.5 V. Battery design is a science and industry in itself, and energy scavenging has attracted a lot of attention in research. This section can only provide some small glimpses of this vast field; some papers that deal with these questions

**Storing energy: Batteries**

**Traditional batteries**

The power source of a sensor node is a battery, either nonrechargeable ("primary batteries") or, if an energy scavenging device is present on the node, also rechargeable ("secondary batteries").

---

**Table 2.2** Energy densities for various primary and secondary battery types [703]

| | | | |
|---|---|---|---|
| | Primary batteries | | |
| Chemistry | Zinc-air | Lithium | Alkaline |
| Energy (J/cm$^3$) | 3780 | 2880 | 1200 |
| | Secondary batteries | | |
| Chemistry | Lithium | NiMHd | NiCd |
| Energy (J/cm$^3$) | 1080 | 860 | 650 |

In some form or other, batteries are electro-chemical stores for energy – the chemicals being the main determining factor of battery technology.
Upon these batteries, very tough requirements are imposed:

**Capacity** They should have high capacity at a small weight, small volume, and low price. The main metric is energy per volume, J/cm$^3$. Table 2.2 shows some typical values of energy densities, using traditional, macroscale battery technologies. In addition, research on "microscale" batteries, for example, deposited directly onto a chip, is currently ongoing.

**Capacity under load** They should withstand various usage patterns as a sensor node can consume quite different levels of power over time and actually draw high current in certain operation modes.

Current numbers on power consumption of WSN nodes vary and are treated in detail in Section 2.2, so it is difficult to provide precise guidelines. But for most technologies, the larger the battery, the more power can be delivered instantaneously. In addition, the rated battery capacity specified by a manufacturer is only valid as long as maximum discharge currents are not exceeded, lest capacity drops or even premature battery failure occurs

**Self-discharge** Their self-discharge should be low; they might also have to last for a long time (using certain technologies, batteries are operational only for a few months, irrespective of whether power is drawn from them or not).

Zinc-air batteries, for example, have only a very short lifetime (on the order of weeks), which offsets their attractively high energy density.

**Efficient recharging** Re charging should be efficient even at low and intermittently available recharge power; consequently, the battery should also not exhibit any “memory effect”.

Some of the energy-scavenging techniques described below are only able to produce cur-rent in the $\mu$A region (but possibly sustained) at only a few volts at best. Current battery technology would basically not recharge at such values.

**Relaxation** Their relaxation effect – the seeming self-recharging of an empty or almost empty battery when no current is drawn from it, based on chemical diffusion processes within the cell – should be clearly understood. Battery lifetime and usable capacity is considerably extended if this effect is leveraged. As but one example, it is possible to use multiple batteries in parallel and “schedule” the discharge from one battery to another, depending on relaxation properties and power requirements of the operations to be supported

**Unconventional energy stores**

Apart from traditional batteries, there are also other forms of energy reservoirs that can be contemplated. In a wider sense, fuel cells also qualify as an electro-chemical storage of energy, directly producing electrical energy by oxidizing hydrogen or hydrocarbon fuels. Fuel cells actually have excellent energy densities (e.g. methanol as a fuel stores 17.6 kJ/cm$^3$), but currently available systems still require a non negligible minimum size for pumps, valves, and so on. A slightly more traditional approach to using energy stored in hydrocarbons is to use miniature versions of heat engines, for example, a turbine . Shrinking such heat engines to the desired sizes still requires a considerable research effort in Micro Electro Mechanical Systems (MEMSs); predictions regarding power vary between 0.1 – 10 W at sizes of about 1 cc . And lastly, even radioactive substances have been proposed as an energy store . Another option are so-called “gold caps”, high-quality and high-capacity capacitors, which can store relatively large amounts of energy, can be easily and quickly recharged, and do not wear out over time.

**DC – DC Conversion**

Unfortunately, batteries (or other forms of energy storage) alone are not sufficient as a direct power source for a sensor node. One typical problem is the reduction of a battery’s voltage as its capacity drops. Consequently, less power is delivered to the sensor node’s circuits, with immediate consequences for oscillator frequencies and transmission power – a node on a weak battery will have a smaller transmission range than one with a full battery, possibly throwing off any calibrations done for the range at full battery ranges.

A DC – DC converter can be used to overcome this problem by regulating the voltage delivered to the node’s circuitry. To ensure a constant voltage even though the battery’s supply voltage drops, the DC – DC converter has to draw increasingly higher current from the battery when the battery is already becoming weak, speeding up battery death (see Figure 3 in reference . Also, the DC – DC converter does consume energy for its own operation, reducing overall efficiency. But the advantages of predictable operation during the entire life cycle can outweigh these disadvantages.

## Energy scavenging

Some of the unconventional energy stores described above – fuel cells, micro heat engines, radioactivity – convert energy from some stored, secondary form into electricity in a less direct and easy to use way than a normal battery would do. The entire energy supply is stored on the node itself – once the fuel supply is exhausted, the node fails.

To ensure truly long-lasting nodes and wireless sensor networks, such a limited energy store is unacceptable. Rather, energy from a node’s environment must be tapped into and made available to the node – **energy scavenging** should take place. Several approaches exist .

**Photovoltaics** The well-known solar cells can be used to power sensor nodes. The available power depends on whether nodes are used outdoors or indoors, and on time of day and whether for outdoor usage. Different technologies are best suited for either outdoor or indoor usage. The resulting power is somewhere between 10 $\mu$W/cm$^2$ indoors and 15 mW/cm$^2$ outdoors. Single cells achieve a fairly stable output voltage of about 0.6 V (and have therefore to be used in series) as long as the drawn current does not exceed a critical threshold, which depends, among other factors, on the light intensity. Hence, solar cells are usually used to recharge secondary batteries. Best trade-offs between complexity of recharging circuitry, solar cell efficiency, and battery lifetime are still open questions.

**Temperature gradients** Differences in temperature can be directly converted to electrical energy. Theoretically, even small difference of, for example, 5 K can produce considerable power, but practical devices fall very short of theoretical upper limits (given by the Carnot efficiency).

See beck effect-based thermoelectric generators are commonly considered; one example is a generator, which will be commercially available soon, that achieves about 80 $\mu$W/cm$^2$ at about 1 V from a 5 Kelvin temperature difference.[7]

**Vibrations** One almost pervasive form of mechanical energy is vibrations: walls or windows in buildings are resonating with cars or trucks passing in the streets, machinery often has low-frequency vibrations, ventilations also cause it, and so on. The available energy depends on both amplitude and frequency of the vibration and ranges from about 0.1 □W/cm$^3$ up to 10, 000 □W/cm$^3$ for some extreme cases (typical upper limits are lower).

Converting vibrations to electrical energy can be undertaken by various means, based on electromagnetic, electrostatic, or piezoelectric principles. Figure 2.4 shows, as an example, a generator based on a variable capacitor Practical devices of 1 cm$^3$ can produce about

□W/cm$^3$ from 2.25 m/s$^2$, 120 Hz vibration sources, actually sufficient to power simple wireless transmitters

**Pressure variations** Somewhat akin to vibrations, a variation of pressure can also be used as a power source. Such piezoelectric generators are in fact used already. One well-known example is the inclusion of a piezoelectric generator in the heel of a shoe, to generate power as a human walks about . This device can produce, on average, 330 □W/cm$^2$. It is, however, not clear how such technologies can be applied to WSNs.

**Flow of air/liquid** Another often-used power source is the flow of air or liquid in wind mills or turbines. The challenge here is again the miniaturization, but some of the work on millimeter-scale MEMS gas turbines might be reusable. However, this has so far not produced any notable results.

To summarize, Table 2.3 gives an overview of typical values of power and energy densities for different energy sources. The values in this table vary somewhat from those presented above as partially different technologies or environments were assumed; all these numbers can only serve as a general orientation but should always be taken with a grain of salt.

□m

□m

Stationary comb

Anchor

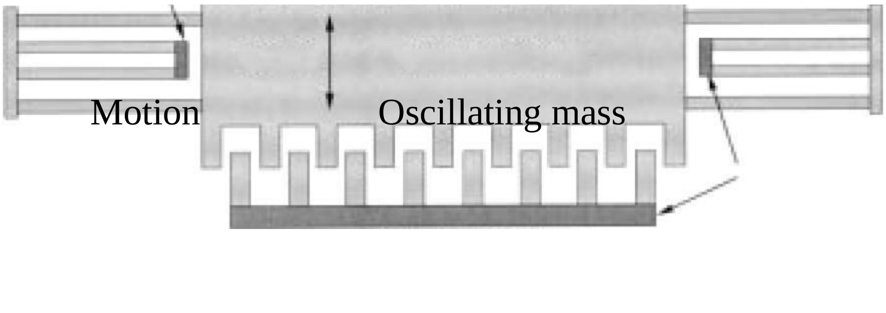

**Figure 2.4** A MEMS device for converting vibrations to electrical energy, based on a variable capacitor [549]. Reproduced by permission of IEEE

**Table 2.3** Comparison of energy sources

| Energy source | Energy density |
|---|---|
| Batteries (zinc-air) | 1050–1560 mWh/$cm^3$ |
| Batteries (rechargeable lithium) | 300 mWh/$cm^3$ (at 3–4 V) |
| **Energy source** | **Power density** |
| Solar (outdoors) | 15 mW/$cm^2$ (direct sun) |
| | 0.15 mW/$cm^2$ (cloudy day) |
| Solar (indoors) | 0.006 mW/$cm^2$ (standard office desk) |
| | 0.57 mW/$cm^2$ (<60 W desk lamp) |
| Vibrations | 0.01–0.1 mW/$cm^3$ |
| Acoustic noise | $3 \cdot 10^{-6}$ mW/$cm^2$ at 75 dB |
| | $9, 6 \cdot 10^{-4}$ mW/$cm^2$ at 100 dB |
| Passive human-powered systems | 1.8 mW (shoe inserts) |
| Nuclear reaction | 80 mW/$cm^3$, $10^6$ mWh/$cm^3$ |