MITIGATION AND EXPLOITATION OF MULTIPATH FADING IN WIRELESS SENSOR NETWORKS

A Thesis

Submitted to the Graduate School
of the University of Notre Dame
in Partial Fulfillment of the Requirements
for the Degree of

Master of Science

by

Daniele Puccinelli, Laurea in Ingegneria Elettronica

Martin Haenggi, Director

Graduate Program in Electrical Engineering
Notre Dame, Indiana
April 2005

CONTENTS

CHAPTER 1

WIRELESS SENSOR NETWORKS

1.1 Overview

The increasing interest in wireless sensor networks can be promptly understood simply by thinking about what they essentially are: a large number of small sensing self-powered nodes which gather information or detect special events and communicate in a wireless fashion, with the end goal of handing their processed data to a base station. Sensing, processing and communication are three key elements whose combination in one tiny device gives rise to a vast number of applications [5, 45]. Sensor networks provide endless opportunities, but at the same time pose formidable challenges, such as the fact that energy is a scarce and usually non-renewable resource. The minimization of energy expenditure in the sensing nodes is a natural direction for wireless sensor network research. Reducing the computational load and minimizing radio power consumption are necessary steps dictated by common sense; on the radio side, the knowledge of the properties of the wireless medium can help minimize packet loss. Wireless analytical models are useful, but the oversimplifications introduced by many of them can lead to wrong considerations and conclusions. For this reason, it is essential to analyze the wireless medium with real hardware. Sensor network design should be performed by considering the interactions between the physical layer and the upper layers: physical layer phenomena should be taken into account in the development of MAC and routing schemes. In the present work,

we consider the impact of multipath fading on sensor networks. Chapter 2 focuses on the analysis of the properties of multipath fading, and the verification of the most common models through software simulations and tests based on hardware implementations. In Chapter 3, we use the properties of multipath fading to propose a novel, lightweight protocol for the acquisition of data from the nodes in a wireless sensor network based on the exploitation of mobility and the concept of opportunistic transmissions. Our protocol aims at the avoidance of packet loss by the nodes and the minimization of their energy consumption.

While Chapter 3 revolves around the mitigation of the effects of multipath fading, chapter 4 describes a novel idea for its constructive exploitation. We show how the footprint of multipath fading on the received signal strength can be used for purposes of motion detection. The present work is completed by an appendix where extra details specific to our hardware implementations are provided to the interested reader.

In the rest of the present chapter, we provide an overview of the research efforts in the field of wireless sensor networks. We start by listing some interesting applications and research projects currently underway. We then indicate the main peculiarities of sensor networks to give a feeling and basic understanding of what makes them unique and to justify the enormous interest in the field that is being displayed by the community. Finally, we present the most interesting and commonly used hardware platforms that enable wireless sensor network research.

1.2 Applications

Possible applications of sensor networks are of interest to the most diverse fields. Environmental monitoring, warfare, child education, terrorist surveillance, microsurgery, and agriculture are only a few examples [24]. Through joint efforts of the

University of California at Berkeley and the College of the Atlantic, environmental monitoring is carried out off the coast of Maine on Great Duck Island by means of a network of *Berkeley motes* equipped with various sensors (temperature, humidity, solar radiation, pressure...) [32]. The nodes send their data to a base station which makes them available on the Internet. Since habitat monitoring is rather sensitive to human presence, the deployment of a sensor network provides a non-invasive approach and a remarkable degree of granularity in data acquisition [12]. The same idea lies behind the Pods project at the University of Hawaii at Manoa [10], where the environmental data (air temperature, light, wind, relative humidity and rainfall) are gathered by a network of weather sensors embedded in the communication units deployed in the South-West Rift Zone in Volcanoes National Park on the Big Island of Hawaii. A major concern of the researchers was in this case camouflaging the sensors to make them invisible to curious tourists. In Princeton's Zebranet Project [3], a dynamic sensor network has been created by attaching special collars equipped with a low-power GPS system to the necks of zebras to monitor their moves and their behavior. Since the network is designed to operate in an infrastructure-free environment, peer-to-peer swaps of information are used to produce redundant databases so that researchers only have to encounter a few zebras in order to collect data. Sensor networks can also be used to monitor and study natural phenomena which intrinsically discourage human presence, such as hurricanes and forest fires. Joint efforts between Harvard University, the University of New Hampshire, and the University of North Carolina have recently led to the deployment of a wireless sensor network to monitor eruptions at Volcán Tungurahua, an active volcano in central Ecuador. A network of Berkeley motes monitored infrasonic signals during eruptions, and data were transmitted over a 9km wireless link to a base station at the volcano observatory [58].

Intel's Wireless Vineyard [11] is an example of using ubiquitous computing for agricultural monitoring. In this application, the network is expected not only to collect and interpret data, but also to use such data to make decisions aimed at detecting the presence of parasites and enabling the use of the appropriate kind of insecticide. Data collection relies on *data mules*, small devices carried by people (or dogs) that communicate with the nodes and collect data. In this project, the attention is shifted from reliable information collection to active decision-making based on acquired data.

Just as they can be used to monitor nature, sensor networks can likewise be used to monitor human behavior. In the Smart Kindergarten project at UCLA [53], wirelessly-networked, sensor-enhanced toys and other classroom objects supervise the learning process of children and allow unobtrusive monitoring by the teacher.

Medical research and healthcare can greatly benefit from sensor networks: vital sign monitoring and accident recognition are the most natural applications. An important issue is the care of the elderly, especially if they are affected by cognitive decline: a network of sensors and actuators could monitor them and even assist them in their daily routine. Smart appliances could help them organize their lives by reminding them of their meals and medications. Sensors can be used to capture vital signs from patients in real-time and relay the data to handheld computers carried by medical personnel, and wearable sensor nodes can store patient data such as identification, history, and treatments. With these ideas in mind, Harvard University is cooperating with the School of Medicine at Boston University to develop CodeBlue, an infrastructure designed to support wireless medical sensors, PDAs, PCs, and other devices that may be used to monitor and treat patients in various medical scenarios [16]. On the hardware side, the research team has created Vital

Dust, a set of devices based on the Mica2[1] sensor node platform (one of the most popular members of the Berkeley motes family), which collect heart rate, oxygen saturation, and EKG data and relay them over a medium-range (100m) wireless network to a PDA [38]. Interactions between sensor networks and humans are already judged controversial. The US has recently approved the use of a radio-frequency implantable device (VeriChip) on humans, whose intended application is accessing the medical records of a patient in an emergency. Potential future repercussions of this decision have been discussed in the media.

An interesting application to civil engineering is the idea of Smart Buildings: wireless sensor and actuator networks integrated within buildings could allow distributed monitoring and control, improving living conditions and reducing the energy consumption, for instance by controlling temperature and air flow. It is estimated that in the USA alone $55 billion a year could be saved, and carbon emissions could be reduced by 35 million metric tons [4].

Military applications are plentiful. An intriguing example is DARPA's self-healing minefield [1], a self-organizing sensor network where peer-to-peer communication between anti-tank mines is used to respond to attacks and redistribute the mines in order to heal breaches, complicating the progress of enemy troops. Urban warfare is another application that distributed sensing lends itself to. An ensemble of nodes could be deployed in a urban landscape to detect chemical attacks, or track enemy movements. PinPtr is an ad hoc acoustic sensor network for sniper localization developed at Vanderbilt University [2, 46]. The network detects the muzzle blast and the acoustic shock wave that originate from the sound of gunfire. The arrival times of the acoustic events at different sensor nodes are used to estimate the position of the sniper and send it to the base station with a special data aggregation

[1]See Section 1.4 for a hardware overview.

and routing service.

Going back to peaceful applications, efforts are underway at Carnegie Mellon University and Intel for the design of IrisNet (Internet-scale Resource-Intensive Sensor Network Services) [18], an architecture for a worldwide sensor web based on common computing hardware such as Internet-connected PCs and low-cost sensing hardware such as webcams. The network interface of a PC indeed senses the virtual environment of a LAN or the Internet rather than a physical environment; with an architecture based on the concept of a distributed database [19], this hardware can be orchestrated into a global sensor system that responds to queries from users.

1.3 Main Features

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

Lifetime. Lifetime is extremely critical for most applications, and its primary limiting factor is the energy consumption of the nodes, which need to be self-powering. Although it is often assumed that the transmit power associated with packet trans-

mission accounts for the lion's share of power consumption, sensing, signal processing and even hardware operation in standby mode consume a consistent amount of power as well [20, 62]. In some applications, extra power is needed for macro-scale actuation.

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

At the physical layer, benefits can be obtained with lower radio duty cycles and dynamic modulation scaling (varying the constellation size to minimize energy expenditure [49]). Using low-power modi for the processor or disabling the radio is generally advantageous, even though periodically turning a subsystem on and off may be more costly than always keeping it on. Techniques aimed at reducing the idle mode leakage current in CMOS-based processors are also noteworthy [14].

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

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

A mobile base station with no energy constraints can decrease energy expenditure at the sensor nodes by reducing the number of necessary transmissions, and eliminating the need for multihop routing and its many undesired effects [27, 25].

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

Untethered systems need to self-configure and adapt to different conditions. Sensor networks should also be robust to changes in their topology, for instance due to the failure of individual nodes. In particular, coverage and connectivity should always be guaranteed. Coverage can be seen as a measure of quality of service in a sensor network [36], as it defines how well a particular area can be observed by a network and characterizes the probability of detection of geographically constrained phenomena or events. Complete coverage is particularly important for surveillance applications. Connectivity refers to the possibility of reaching the base station from any node.

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

the normal operation of the node [43].

The functioning of the network as a whole should not be endangered by unavoidable failures of single nodes, which may occur for a number of reasons, from battery depletion to unpredictable external events, and may either be independent or geographically correlated [17]. Fault-tolerance is particularly crucial as ongoing maintenance is rarely an option in sensor network applications.

Self-configuring nodes are necessary to allow the deployment process to run smoothly without human interaction, which should in principle be limited to placing nodes into a given geographical area. It is not desirable to have humans configure nodes for habitat monitoring and destructively interfere with wildlife in the process, or configure nodes for urban warfare monitoring in a hostile environment. The nodes should be able to assess the quality of the network deployment and indicate any problems that may arise, as well as adjust to changing environmental conditions by automatic reconfiguration. Location awareness is important for self-configuration and has definite advantages in terms of routing [34] and security. Time synchronization [52] is advantageous in promoting cooperation among nodes, such as data fusion, channel access, coordination of sleep modi, or security-related interaction.

Data Collection. Data collection is related to network connectivity and coverage. An interesting solution is the use of ubiquitous mobile agents that randomly move around to gather data bridging sensor nodes and access points, whimsically named data MULEs (Mobile Ubiquitous LAN Extensions) in [50]. The predictable mobility of the data sink can be used to save power [13], as nodes can learn its schedule. A similar concept has been implemented in Intel's Wireless Vineyard.

It is often the case that all data are relayed to a base station, but this form of centralized data collection may shorten network lifetime. Relaying data to a data

sink causes non-uniform power consumption patterns that overburden forwarding nodes [25]. This is particularly harsh on nodes providing end links to base stations, which may end up relaying traffic coming from all other nodes, thus forming a critical bottleneck for network throughput [24, 22], as shown in Figure 3.3.

An interesting technique is clustering [61]: nodes team up to form clusters and transmit their information to their cluster heads, which fuse the data and forward it to a sink. Fewer packets are transmitted, and a uniform energy consumption pattern may be achieved by periodic re-clustering. Data redundancy is minimized, as the aggregation process fuses strongly correlated measurements.

Many applications require that queries be sent to sensing nodes. This is true, for example, whenever the goal is gathering data regarding a particular area where various sensors have been deployed. This is the rationale behind looking at a sensor network as a database [21].

A sensor network should be able to protect itself and its data from external attacks, but the severe limitations of lower-end sensor node hardware make security a true challenge. Typical encryption schemes, for instance, require large amounts of memory that are unavailable in sensor nodes. Data confidentiality should be preserved by encrypting data with a secret key shared with the intended receiver. Data integrity should be ensured to prevent unauthorized data alteration. An authenticated broadcast must allow the verification of the legitimacy of data and their sender. In a number of commercial applications, a serious disservice to the user of a sensor network is compromising data availability (denial of service), which can be achieved by sleep-deprivation torture [54]: batteries may be drained by continuous service requests or demands for legitimate but intensive tasks [33], preventing the node from entering sleep modi.

1.4 Available Hardware Platforms

Berkeley motes, made commercially available by Crossbow, are by all means the best known sensor node hardware implementation, used by more than 100 research organizations. They consist of an embedded microcontroller, low-power radio, and a small memory, and they are powered by two AA batteries. At this point, they represent a very good solution for generic sensing nodes, even though their unit cost is still relatively high (about $200). Mica and Mica2 are the most successful families of Berkeley motes. The Mica2 platform is equipped with an Atmel ATmega128L and has a CC1000 transceiver. A 51-pin expansion connector is available to interface sensors (commercial sensor boards designed for this specific platform are available). Since the MCU is to handle medium access and baseband processing, a fine-grained event-driven real-time operating system (TinyOS) has been implemented to specifically address the concurrency and resource management needs of sensor nodes. For applications that require a better form factor, the circular Mica2Dot can be used: it has most of the resources of Mica2, but is only 2.5cm in diameter. Berkeley motes up to the Mica2 generation cannot interface with other wireless-enabled devices [9]. However, the newer generations Mica-Z and Telos support IEEE 802.15.4, which is part of the 802.15 Wireless Personal Area Network (WPAN) standard being developed by IEEE.

Intel has designed its own iMote [28] to implement various improvements over available mote designs, such as increased CPU processing power, increased main memory size for on-board computing and improved radio reliability. In the iMote, a powerful ARM7TDMI core is complemented by a large main memory and non-volatile storage area; on the radio side, Bluetooth has been chosen.

Various platforms have been developed for the use of Berkeley motes in mobile sensor networks to enable investigations into controlled mobility, which facilitates

deployment and network repair (like in the self-healing minefields) and provides possibilities for the implementation of energy-harvesting. UCLA's RoboMote [15], Notre Dame's MicaBot [35] and UC Berkeley's CotsBots [7] are examples of efforts in this direction.

UCLA's Medusa MK-2 sensor nodes [47], developed for the Smart Kindergarten project, expand Berkeley motes with a second microcontroller. An on-board power management and tracking unit monitors power consumption at the different subsystems and selectively powers down unused parts of the node. The Medusa MK-2 nodes can also act as gateways. A MEMS accelerometer and a temperature sensor form an on-board sensing subsystem, and other sensors on expansion boards can be added.

UCLA has also developed iBadge [40], a wearable sensor node with sufficient computational power to process the sensed data. Built around an ATMega128L and a DSP, it features a Localization Unit designed to estimate the position of iBadge in a room based on the presence of special nodes of known location attached to the ceilings.

In the context of the EYES project (a joint effort among several European institutions) custom nodes [56, 43] have been developed to test and demonstrate energy-efficient networking algorithms. On the software side, a proprietary operating system, PEEROS (Preemptive EYES Real Time Operating System), has been implemented.

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

The DIY Smart-its [55] have been developed in the UK at Lancaster University; their modular design is based on a core board that provides processing and com-

munication and can be extended with add-on boards. A typical setup of Smart-its consists of one or more sensing nodes that broadcast their data to a base station which consists of a standard core board connected to the serial port of a PC. Simplicity and extensibility are the key features of this platform, which has been developed for the creation of Smart Objects. An interesting application is the Weight Table: four load cells placed underneath a coffee table form a Wheatstone bridge and are connected to a DIY node that observes load changes, determines event types like placement and removal of objects or a person moving a finger across the surface, and also retrieves the position of an object by correlating the values of the individual load cells after the event type (removed or placed) has been recognized [48].

Particle Computers have been developed at the University of Karlsruhe, Germany. Similarly to the DIY platform, the Particle Smart-its are based on a core board equipped with a Microchip PIC; they are optimized for energy efficiency, scalable communication and small scale (17mm $\times$ 30mm). Particles are able to communicate in an ad hoc fashion: if two Particles come close to one another, they can instantly talk to each other. Additionally, if Particles come near a gateway device, they can be connected to Internet-enabled devices and access services and information on the Internet as well as provide information [6].

The BTnode hardware from ETHZ [9] is based on an Atmel ATmega128L microcontroller and a Bluetooth module. Although advertised as a low-power technology, Bluetooth has a relatively high power consumption, as discussed before. It also has long connection setup times and a lower degree of freedom with respect to possible network topologies. On the other hand, it ensures interoperability between different devices, enables application development through a standardized interface, and offers a significantly higher bandwidth (about 1 Mbps) compared to many low-power radios (about 50 Kbps). Moreover, Bluetooth support means that COTS hardware

can be used to create a gateway between a sensor network and an external network (*e.g.*, the Internet), as opposed to more costly proprietary solutions [8].

MIT is working on the μAMPS (μ-Adaptive Multi-domain Power-aware Sensors) project, which explores energy-efficiency constraints and key issues such as self-configuration, reconfigurability, and flexibility. A first prototype has been designed with COTS components: three stackable boards (processing, radio and power) and an optional extension module. The energy dissipation of this microsensor node is reduced through a variety of power-aware design techniques [37] including fine-grain shutdown of inactive components, dynamic voltage and frequency scaling of the processor core, and adjustable radio transmission power based on the required range. Dynamic voltage scaling is a technique used for active power management where the supply voltage and clock frequency of the processor are regulated depending on the computational load, which can vary significantly based on the operational mode [14, 62]. The main goal of second generation μAMPS is clearly stated in [57] as breaking the 100 μW average power barrier.

Another interesting MIT project is the Pushpin computing system [31], whose goal is the modeling, testing, and deployment of distributed peer-to-peer sensor networks consisting of many identical nodes. The pushpins are 18mm $\times$ 18mm modular devices with a power substrate, an infrared communication module, a processing module (Cygnal C8051F016) and an expansion module (*e.g.*, for sensors); they are powered by direct contact between the power substrate and layered conductive sheets.

MIT has also built Tribble (Tactile reactive interface built by linked elements), a spherical robot wrapped by a wired skinlike sensor network designed to emulate the functionalities of biological skin [39]. Tribble's surface is divided into 32 patches with a Pushpin processing module and an array of sensors and actuators.

At Lancaster University, surfaces provide power and network connectivity in the Pin&Play project. Network nodes come in different form factors, but all share the Pin&Play connector, a custom component that allows physical connection and networking through conductive sheets which are embedded in surfaces such as a wall or a bulletin board [29]. Pin&Play falls in between wired and wireless technologies as it provides network access and power across 2D surfaces. Wall-mounted objects are especially suited to be augmented to become Pin&Play objects. In a demonstration, a wall switch was augmented and freely placed anywhere on a wall with a Pin&Play surface as wallpaper.

For applications which do not call for the minimization of power consumption, high-end nodes are available. Rockwells WINS nodes and Sensoria's WINS 3.0 Wireless Sensing Platform are equipped with more powerful processors and radio systems. The embedded PC modules based on widely-supported standards PC/104 and PC/104-plus feature Pentium processors; moreover, PC/104 peripherals include digital I/O devices, sensors and actuators, and PC-104 products support almost all PC software. PFU Systems's Plug-N-Run products, which feature Pentium processors, also belong to this category. They offer the capabilities of PCs and the size of a sensor node, but lack built-in communication hardware. COTS components or lower-end nodes may be used in this sense [60]. Research is underway toward the creation of sensor nodes that are more capable than the motes, yet smaller and more power-efficient than higher-end nodes.

Simple yet effective gateway devices are the MIB programming boards from Crossbow, which bridge networks of Berkeley motes with a PC (to which they interface using the serial port or Ethernet). In the case of Telos motes, any generic node (*i.e.*, any Telos Mote) can act as a gateway, as it may be connected to the USB port of a PC and bridge it to the network. Of course, more powerful gateway

devices are also available. Crossbow's Stargate is a powerful embedded computing platform (running Linux) with enhanced communication and sensor signal processing capabilities based on Intel PXA255, the same X-Scale processor that forms the core of Sensoria WINS 3.0 nodes. Stargate has a connector for Berkeley motes, may be bridged to a PC via Ethernet or 802.11, and includes built-in Bluetooth support.

CHAPTER 2

STATIC MULTIPATH FADING

2.1 Fading Models

Multipath fading [42, 26, 44] is typically observed in rich scattering environments, such as indoor locales where walls and all sorts of objects cause reflection of the radio-frequency signals. The literature usually considers multipath fading in a dynamic context: the receiver is assumed to be mobile.

Radio waves travel through free space undergoing various phenomena such as absorption, reflection, refraction, diffraction, and scattering. The atmosphere, the topology of the environment and the objects in their path are responsible for most of the characteristic features of the received signal. There may or may not exist a line-of-sight (LOS) path between the transmitter and the receiver. In such cases, propagation is obtained by means of reflection and scattering from the physical obstacles and diffraction over or around them. The transmitted signal reaches the receiver via several paths with different time delays; hence the term *multipath*. At the receiver, these multipath waves with different amplitudes and phases combine to give a resultant signal whose mean strength is given by the large-scale path loss; the deviation from the mean is referred to as *fading*.

Multipath fading is a *small-scale* phenomenon: the level of attenuation of the signal that it causes changes substantially if the position of the receiver is varied by about half a wavelength. Shadowing is a *large-scale* effect, as it models substantial

deviations from the mean due to large obstacles. Multipath fading can inflict several fades upon a moving receiver in a short interval. Large obstacles create shadow zones which cause deep fades if a receiver happens to enter them.

Small-scale fading can be further classified as flat or frequency selective, and slow or fast. A received signal is said to undergo *flat fading* if the mobile radio channel has a constant gain and a linear phase response over a bandwidth greater than the one of the transmitted signal. The received signal is characterized by amplitude fluctuations as a result of the variations in the channel gain caused by the multipath effects, but the shape of its spectrum remains intact at the receiver. If the radio channel has a constant gain and linear phase response over a bandwidth smaller than that of the transmitted signal, the latter is said to undergo *frequency selective fading*. The received signal consists of multiple versions of the transmitted signal which are attenuated and delayed. The result is time dispersion of the transmitted symbols within the channel arising from these different time delays causing intersymbol interference (ISI). When there is relative motion between the transmitter and the receiver, Doppler spread is introduced in the received signal spectrum causing frequency dispersion. If the Doppler spread is significant with respect to the bandwidth of the transmitted signal, the received signal is said to undergo *fast fading*. The impulse response of the channel changes rapidly within a symbol duration, causing the coherence time of the channel to be smaller than a symbol period. However, if the Doppler spread of the channel is much less than the bandwidth of the baseband signal, the signal is said to undergo *slow fading*.

One of the most common features of wireless sensor networks is the fact that the nodes are static: this is the reason why this chapter focuses on static multipath fading. For sensor network, small-scale fading effects (in particular, multipath fading) are dominant. In particular, flat fading occurs with nodes using narrowband radios.

Fast fading is usually not an issue: the fading level is approximately constant over a symbol period (in the order of tens of μs). Large-scale fading is also of interest in the form of shadowing.

2.1.1 Rayleigh Fading

If we place a receiver in a fixed position in a rich scattering environment, the received signal that it measures is the result of the superposition of all the scattered paths that reach it. If we assume the transmitted signal to be vertically polarized, the electric field measured at the transmitter is given by the sum of the N rays hitting the receiver:

$$E_z = E_0 \sum_{n=1}^{N} C_n \cos(\omega_c t + \phi_n), \tag{2.1}$$

where ω_c is the carrier frequency of the transmitted signal, and $E_0 C_n$ indicates the amplitude of the nth wave contributing to E_z. We assume that none of these components dominates over the others (lack of a direct line-of-sight path). If we define the two random processes describing the in-phase and quadrature components of E_z, respectively, as

$$T_c = E_0 \sum_{n=1}^{N} C_n \cos(\omega_c t + \phi_n) \tag{2.2}$$

and

$$T_s = E_0 \sum_{n=1}^{N} C_n \sin(\omega_c t + \phi_n), \tag{2.3}$$

we can express E_z as

$$E_z = T_c(t) \cos(\omega_c t + \phi_n) - T_s(t) \sin(\omega_c t + \phi_n). \tag{2.4}$$

For large values of N, by the Central Limit Theorem, T_c and T_s are Gaussian random processes. For a fixed t, T_c and T_s are zero-mean Gaussian random variables with equal variance $\sigma^2 = \frac{E_0^2}{2}$. The mean and the variance are obtained by integrating over all the incidence angles α_n. T_c and T_s are uncorrelated, and therefore independent. The envelope of E_z is given by

$$R = \sqrt{T_c^2 + T_s^2} \tag{2.5}$$

and is Rayleigh distributed:

$$p_R(r) = \begin{cases} \frac{r}{\sigma^2} \exp\left(-\frac{r^2}{2\sigma^2}\right) & \text{if } r \geq 0 \\ 0 & \text{if } r < 0 \end{cases}$$

Since the sum of the squares of two Gaussian random variables is an exponential random variable, the density of the squared envelope R^2 is exponential.

2.1.2 Ricean Fading

In the presence of a dominant, static line-of-sight component, the envelope obeys a Ricean distribution. Its probability density may be written as

$$p(r) = \begin{cases} \frac{r}{\sigma^2} \exp\left(-\frac{(r^2+A^2)}{2\sigma^2}\right) I_0\left(\frac{Ar}{\sigma^2}\right) & \text{if } r \geq 0, \text{ with } A \geq 0 \\ 0 & \text{if } r < 0 \end{cases}$$

The parameter A indicates the peak amplitude of the dominant path, and I_0 is the modified Bessel function of the first kind and zero-order. The Ricean factor K is defined as

$$K(\mathbf{dB}) = 10 \log \frac{A^2}{2\sigma^2} \tag{2.6}$$

and is measured in dB. As $A \to 0$ and $K \to -\infty$ the Ricean distribution degenerates to a Rayleigh distribution (the dominant path is no longer present).

2.2 Multipath Static Fading

2.2.1 The Fading Function

The Rayleigh and Ricean models consider fading from a probabilistic point of view. In the presence of a line of sight, if we place a receiver in different positions at the same distance with respect to a transmitter and measure the received signal strength, we expect to obtain a collection of samples following a Ricean distribution. By the same token, in the absence of a line of sight, our samples are expected to follow a Rayleigh distribution. This point of view often casts a shadow on the fact that fading is indeed a deterministic phenomenon. If we place a transmitter and a receiver in a room and allow them to communicate, the received signal strength is uniquely determined by the topology of the room, and does not fluctuate in time. Time fluctuations come into the picture if the terminals move around. For instance, a mobile receiver will see a different topology depending on its position with respect to the transmitter.

Let us consider a wireless point-to-point setting, which can be generalized to a network. A transmitter T and a receiver R are deployed in a rich scattering environment. The position of the pair of nodes is uniquely determined by the two 3-tuples indicating their spatial coordinates with respect to a fixed reference, namely $(x_{\mathrm{T}}, y_{\mathrm{T}}, z_{\mathrm{T}})$ and $(x_{\mathrm{R}}, y_{\mathrm{R}}, z_{\mathrm{R}})$. If we assume a multipath channel with additive white Gaussian noise, the signal from T received by R at time k is

$$y_k = a_k x_k + z_k, \tag{2.7}$$

where a_k is the large-scale path loss multiplied by the fading coefficient, whose inverse we shall refer to as the (instantaneous) channel gain. Multipath fading depends on the topology of the environment which causes the signal from T to reach R by means of different paths depending on the position of the pair with respect to

the topology of the environment. The large-scale path loss and the net effect of the multiple paths determine a channel gain which can be obtained experimentally by averaging y_k over time to eliminate the noise z_k.

Let us define a *fading function* $f : \mathbb{R}^6 \rightarrow \mathbb{R}^+$ which deterministically maps the spatial coordinates of the transmitter and the receiver to a coefficient, modeling the effect of multipath fading. In a network, the fading function associates a channel gain to the links between each node pair. Multipath fading favors some channels and penalizes others as a function of the position of the nodes with respect to the environment; this is particularly relevant if the nodes are static (lack of spatial diversity). It should be noted that the fading function does not depend on time, and that it does not incorporate the effects of the large scale path loss L. Hence, the channel gain may be written as

$$a_k = L \cdot f(x_{\mathrm{T}}, y_{\mathrm{T}}, z_{\mathrm{T}}, x_{\mathrm{R}}, y_{\mathrm{R}}, z_{\mathrm{R}}). \tag{2.8}$$

We will provide a clear illustration of these concepts by means of simulation as well as experimental results. We have created a C simulator that implements a simple model to extrapolate the fading function for a two-dimensional setting where T occupies a fixed position within a rectangular room and R is placed at a number of different positions; in other words, our simulator computes a planar fading function $f_{2\mathrm{D}} : \mathbb{R}^4 \rightarrow \mathbb{R}^+$ mapping $(x_{\mathrm{T}}, y_{\mathrm{T}}, x_{\mathrm{R}}, y_{\mathrm{R}})$ to a fading coefficient.

The pair $(x_{\mathrm{T}}, y_{\mathrm{T}})$ is fixed and $(x_{\mathrm{R}}, y_{\mathrm{R}})$ is allowed to vary in order to cover the two-dimensional topology. The output of the simulator is a mapping of the room where a signal strength level is associated with each point for a given position of the transmitter. The simulator traces N paths originating from T to check whether they reach R with reasonable strength (at least -110dBm, which is a typical value for the kind of radios used in sensor node platforms). A simple path loss model such

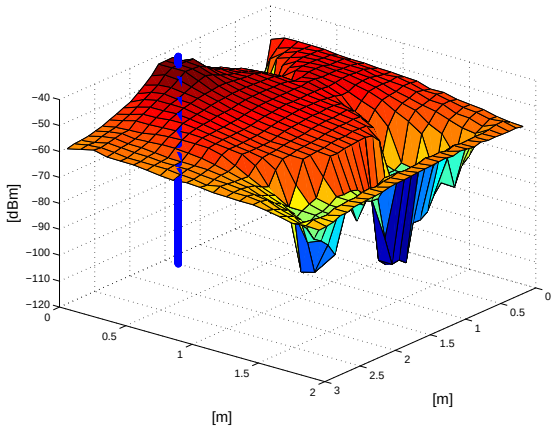

Figure 2.1. The received power for a rectangular topology with the transmitter in the position (1.5m, 0.1m) as computed by our simulator. The position of the transmitter is indicated. The product of the fading function and the large-scale path loss is simply a shifted version of the received power, as the transmitted power is constant.

as

$$P_{rec} = P_0 d^{\alpha} \tag{2.9}$$

is used to account for the large-scale path loss. The power at a reference distance of 1m P_0 has been set to -50dBm on the basis of our experimental experience, and α has been chosen to be 1.8 in accordance with [42] and our experimental results. Near-field effects are not taken into account. A Wall Attenuation Factor of 10dB has been chosen to account for the additional loss due to the fact that the signals are reflected off the walls. Phase changes have been accounted for by measuring the total path from T to R and looking at the remainder of the division between the path and the wavelength; an additional phase change of π is factored in for each reflection. As for the geometrical tracing of the rays, it is assumed that the incidence and reflection angle are equal.

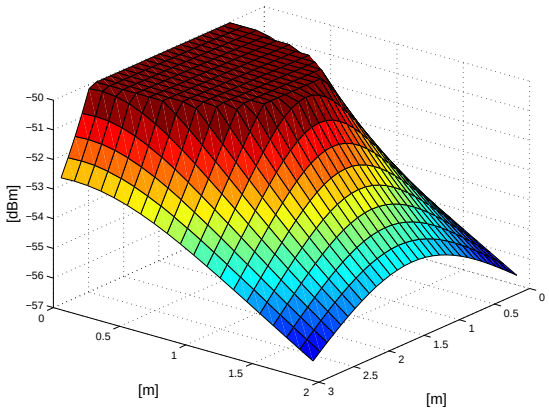

Figure 2.2. The path loss mapping for a rectangular topology with the transmitter in the position (1.5m, 0.1m) as computed by our simulator, ignoring near-field effects (within 1m from the transmitter) for simplicity.

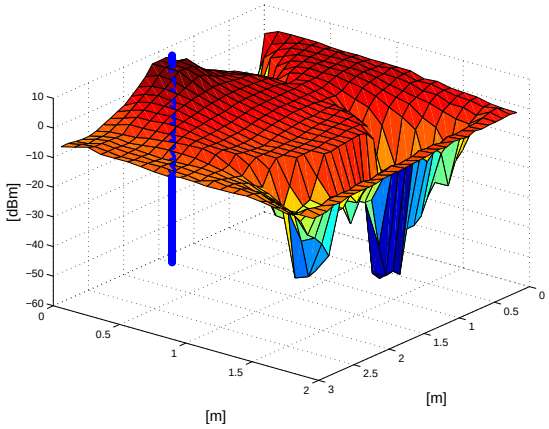

Figure 2.3. The received power without the large-scale path loss for a rectangular topology with the transmitter in the position (1.5m, 0.1m) as computed by our simulator. The position of the transmitter is indicated.

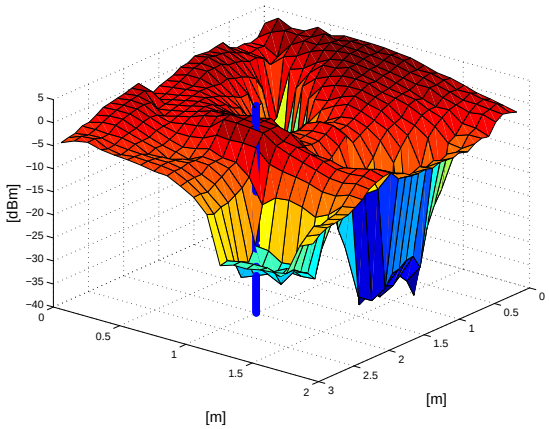

Figure 2.4. The received power without the large-scale path loss for a rectangular topology with the transmitter in the position (2m, 1m) as computed by our simulator. The position of the transmitter is indicated.

Given a two-dimensional orthogonal coordinate system, we define a rectangular topology modeling a room measuring 3m $\times$ 2m. Figure 2.1 shows the received power when the transmitter is located in the position (1.5m, 0.1m). Figure 2.2 shows the path loss mapping for the same topology, ignoring near-field effects. Figure 2.3 shows the received power without the large-scale path loss; comparing this Figure with Figure 2.1 and Figure 2.2 shows that multipath fading dominates over the large-scale path loss. This leads to the observation that the value of the path loss exponent is not that important so long as small distances (smaller than 10m) are considered.

Let us use the simulator to illustrate an important point: the fading function is a mapping from $\mathbb{R}^6$ ($\mathbb{R}^4$ with our two-dimensional simulator) into $\mathbb{R}$. This means that the fading function assigns a different channel to different pairs of transmitters and receivers, which explains why mappings pertaining to the same geometry but with different locations for the transmitter are substantially different. The geometry of

the room is exactly the same, but the position of the transmitter has been modified to (2m, 1m). The result is shown in Figure 2.4: a totally different mapping, to illustrate how the key element is not just the position of the receiver, but its position with respect to the transmitter. Surfaces such as the ones in Figures 2.3 and 2.4 only provide the mapping performed by the fading function for one position of the transmitter; as the transmitter moves, the surface keeps evolving: peaks form and merge back, troughs vanish and reappear.

2.2.2 A deterministic, spatial phenomenon

The two key points that we wish to illustrate are the deterministic and spatial nature of multipath fading. We can use our simulator to measure the signal strength along a circumference centered at some point in the room with the transmitter in a fixed position. We have performed the same experiment with real hardware by placing the receiver on a motorized turntable.

As the receiver occupies different positions along the circumference, the fading function assigns a different channel gain to the transmitter-receiver pair. The channel is expected to remain reasonably constant within displacements smaller than $\frac{\lambda}{2}$, which is indeed the case in the simulation in Figure 2.5 where we have chosen $\frac{\lambda}{2}$=34.5cm (which corresponds to a carrier frequency of 433MHz, the same as in our Mica2 hardware). We expect the waveform shown in Figure 2.5 to be repeated periodically as the receiver keeps moving along the circle, because it keeps occupying the same positions which are mapped to the same channel gain by the deterministic and time-independent fading function. Our simulator would obviously show the periodicity of the received signal strength, but its simplicity calls for the need for experimentation with real hardware to verify our claim.

Figure 2.6 shows the signal strength measured during a point-to-point transmission

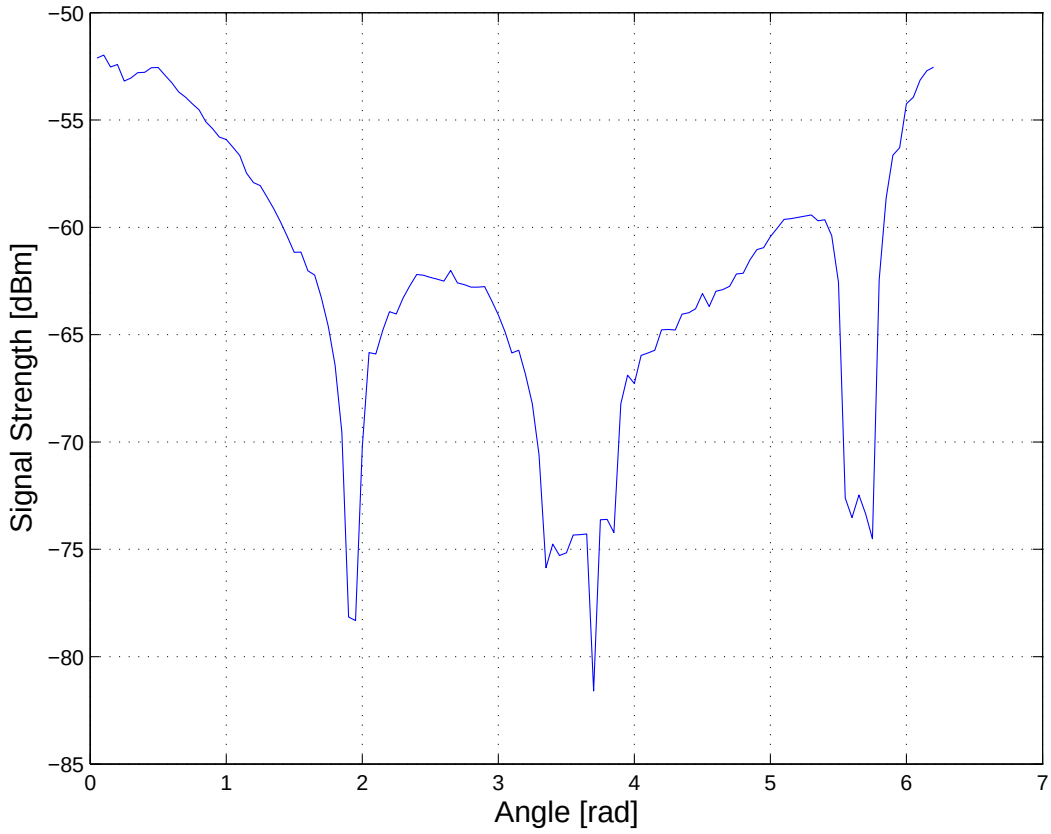

Figure 2.5. Simulation output: received signal strength as measured by a receiver moving along a circle of radius 1m centered at (3m, 2m) in a rectangular environment of 5m $\times$ 3m with a transmitter at (2m, 2m). In this figure, $\frac{\lambda}{2}$ corresponds to 0.345 rad.

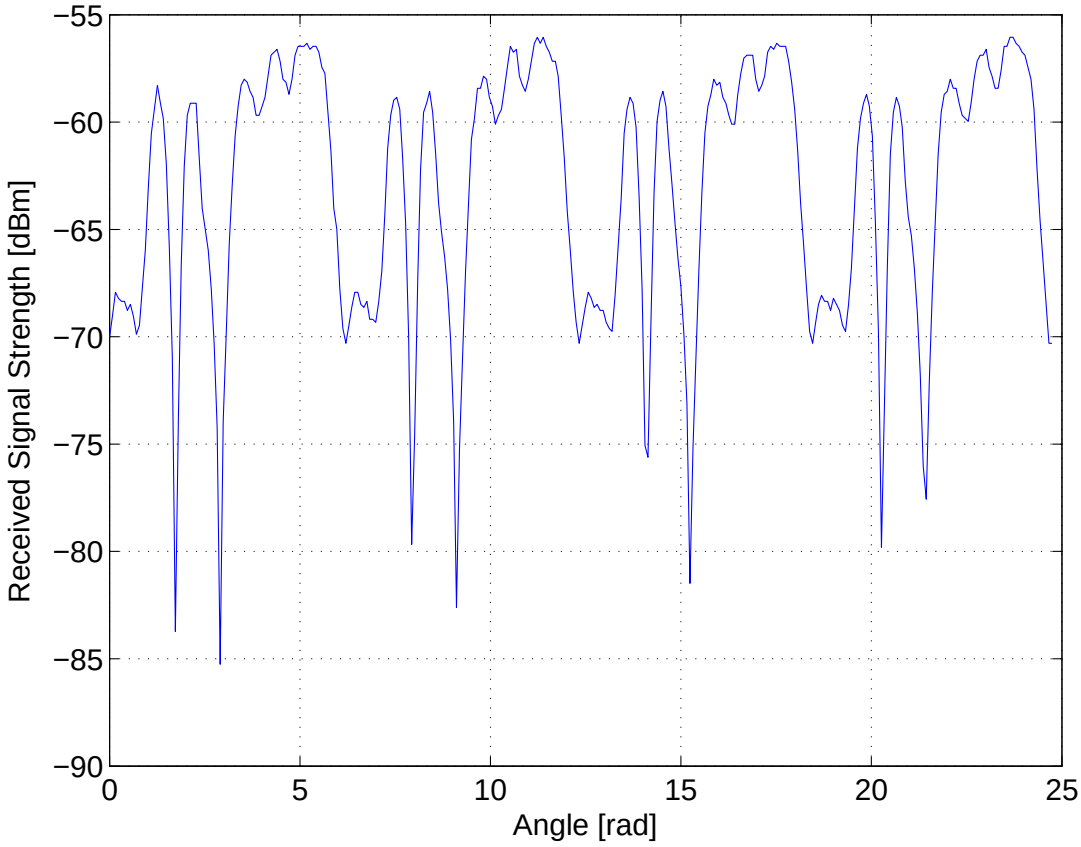

Figure 2.6. Spatial nature of static fading: signal acquisition with Mica2 hardware. The motorized turntable is rotating at 7.5 rpm, and $\frac{\lambda}{2}$ corresponds to 0.57m. In this figure, $\frac{\lambda}{2}$ corresponds to 0.345 rad.

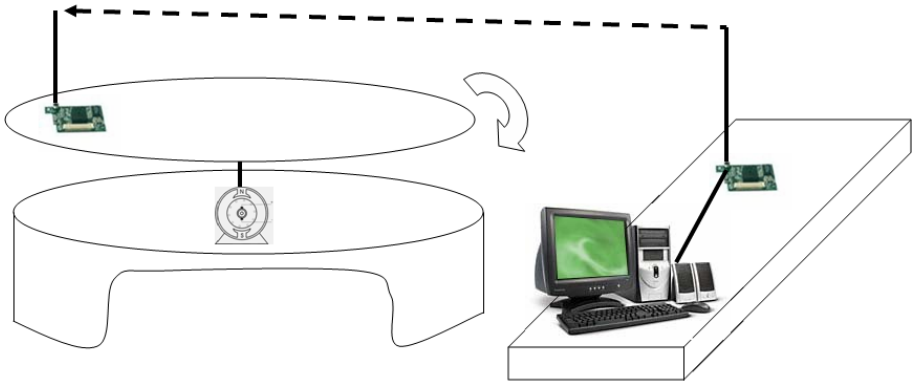

Figure 2.7. Our standard setting for the turntable test.

between Mica2 motes. For this particular experiment, a receiver was placed on a desk and a transmitter was located on a large motorized turntable; a line-of-sight path was maintained between them at all times. We will refer to this setting, portrayed in Figure 2.7 as our *standard turntable test*. From Figure 2.6, it is clear how the particular channel seen by the receiver only depends on the angle of the turntable. The signal shown is periodic for all practical purposes; it does not appear to be exactly periodic owing to quantization and sampling (the period of the revolution of the turntable is not an integer multiple of the sampling time). The level of multipath fading for the link only depends on the position of the pair of nodes and the spatial characterization of the room.

A comparison between Figure 2.5 and 2.6 also shows the validity of our simple simulator: despite its simplicity, it yields very reasonable results.

Since the time variations of the RSSI are uniquely due to the motion of the turntable, the rotational speed of the turntable directly influences the frequency contents of RSSI; this can be clearly seen in the frequency domain in Figures 2.8-2.11. Energy patterns are repeated with a frequency directly depending on the speed of the turntable. This shows the spatial nature of fading from yet another point of view.

It is worth observing that channel gains generally tend to be symmetric: the mapping $f(x_{\mathrm{T}}, y_{\mathrm{T}}, z_{\mathrm{T}}, x_{\mathrm{R}}, y_{\mathrm{R}}, z_{\mathrm{R}})$ tends to be highly correlated with the mapping $f(x_{\mathrm{R}}, y_{\mathrm{R}}, z_{\mathrm{R}}, x_{\mathrm{T}}, y_{\mathrm{T}}, z_{\mathrm{T}})$.

2.2.3 Multipath Fading and Indoor Scenarios

It is often assumed that a wideband radio is more immune to multipath fading than a narrowband radio. If we look at the current status of sensor network research platforms, platforms such as Mica and Mica2 have narrowband radios, whereas Telos

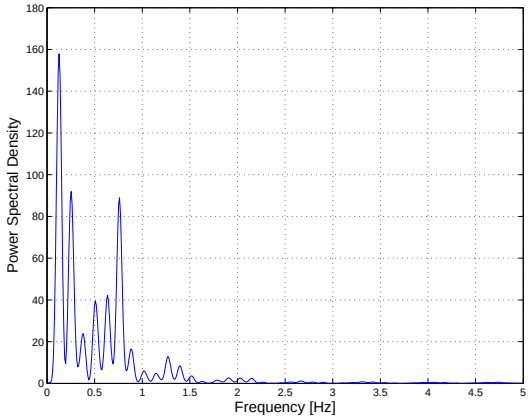

Figure 2.8. At 7.5 rpm, it takes about 8 seconds to complete a revolution; hence power density is higher at integer multiples of $\frac{1}{8} = 125$ mHz.

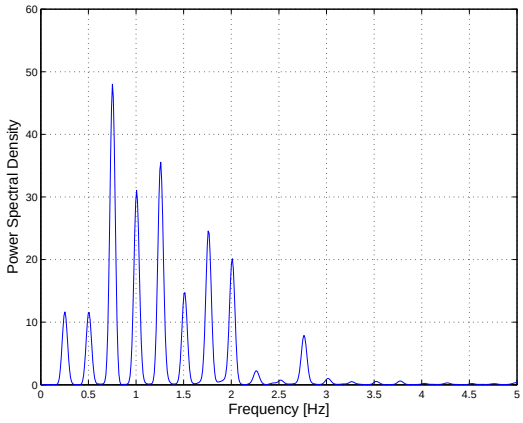

Figure 2.9. At 15.4 rpm, it takes 3.9 seconds to complete a revolution; hence power density is higher at integer multiples of $\frac{1}{3.9} = 256$ mHz.

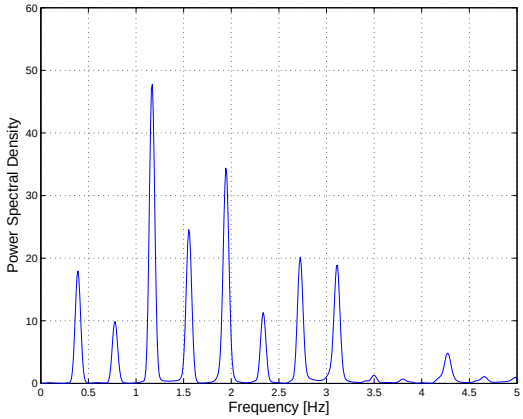

Figure 2.10. At 32 rpm, it takes 2.61 seconds to complete a revolution; hence power density is higher at integer multiples of $\frac{1}{2.61} = 383$ mHz.

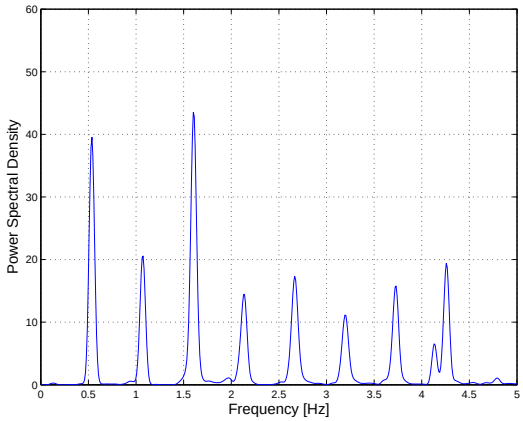

Figure 2.11. At 32 rpm, it takes 1.86 seconds to complete a revolution; hence power density is higher at integer multiples of $\frac{1}{1.86} = 540$ mHz.

and MicaZ have a wideband radio, Chipcon's CC2420. CC2420 has a carrier frequency of 2.4GHz and operates Direct Sequence Spread Spectrum with a spreading gain of 9dB; it is often remarked how solidly it can mitigate multipath fading.

In a wireless channel, the delay spread Δt represents the time it takes the radio signal to cover the path difference Δl. The phase change relative to Δl is given by

$$\frac{\Delta l}{\lambda} = \Delta t \cdot c \cdot \frac{f}{c} = \Delta t \cdot f. \tag{2.10}$$

If $\frac{\Delta l}{\lambda} = 1$, the phase changes by 2π, and we can define the coherence bandwidth as [42]

$$W_c \approx \frac{1}{2\pi\Delta t}, \tag{2.11}$$

where

$$\Delta t = \frac{\Delta l}{\lambda \cdot f} \tag{2.12}$$

The path difference Δl depends on the topology of the environment where the nodes are deployed. For indoor locales such as offices, 10m is a reasonable upper bound; thus a reasonable value for the coherence bandwidth is W_c=5MHz. The CC1000 radio on the Mica2 motes has a radio signal bandwidth of less than 100kHz, which is considerably less than the coherence bandwidth of the channel and causes flat fading. In the case of the CC2420 on MicaZ and Telos, the signal bandwidth happens to be in the same ballpark as the channel coherence bandwidth, and it is to be expected that the use of direct sequence spread spectrum will not help as much as it would with frequency-selective fading, *i.e.*, a situation where the bandwidth of the signal was much wider than W_c. This expectation is indeed confirmed by a simple turntable test performed with a couple of Telos motes. Figure 2.12 shows

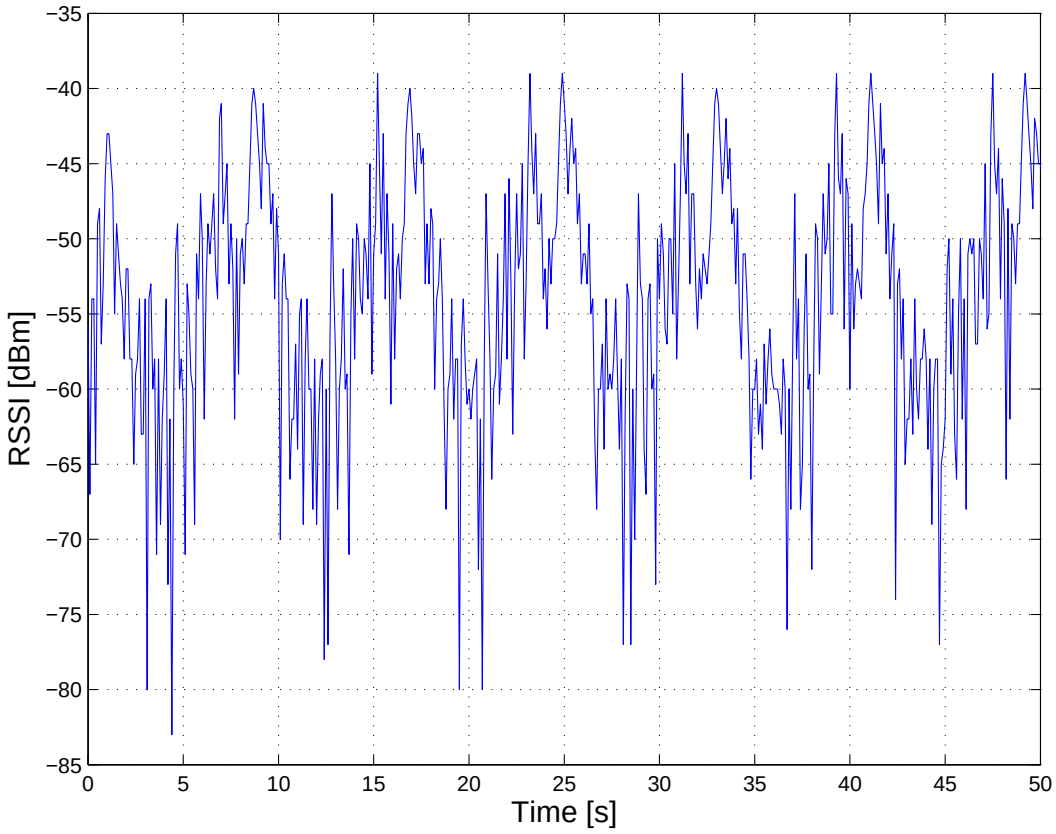

Figure 2.12. Signal strength acquisition with Telos hardware; the turntable rotates at 7.5rpm. Since the carrier frequency is 2.4GHz, $\frac{\lambda}{2}$ is approximately 6cm and corresponds to 0.098rad on this plot.

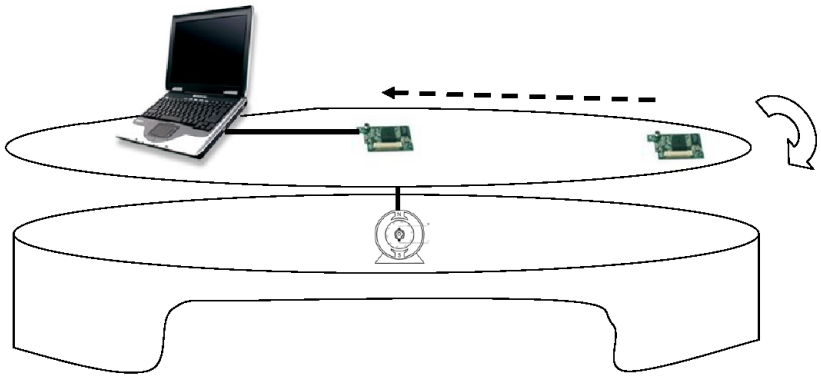

Figure 2.13. In this setting, the dominant component is completely static.

the received signal, affected by multipath fading in a measure perfectly comparable to the case of Mica2. This also explains the difference in the radio range between indoor and outdoor deployments experienced with Telos. In outdoor experiments, a range of around 110m has been measured, and MoteIV claims that only about 50m can be obtained indoor.

2.2.4 The convergence of the probabilistic and deterministic approaches

The probabilistic and deterministic viewpoint are two sides of the same coin. Since computing the fading function is too complex and topology-dependent, it makes sense to work out a probabilistic model. The probabilistic approach provides guidelines as to the distribution of the received signal strength samples. The deterministic angle clarifies that the measured signal strength is not a random variable, but a deterministic quantity set by the topology of the environment. So, if we look at the distribution of our signal strength samples provided by our simulation and our experiment, we expect a Ricean distribution, as a line-of-sight component was

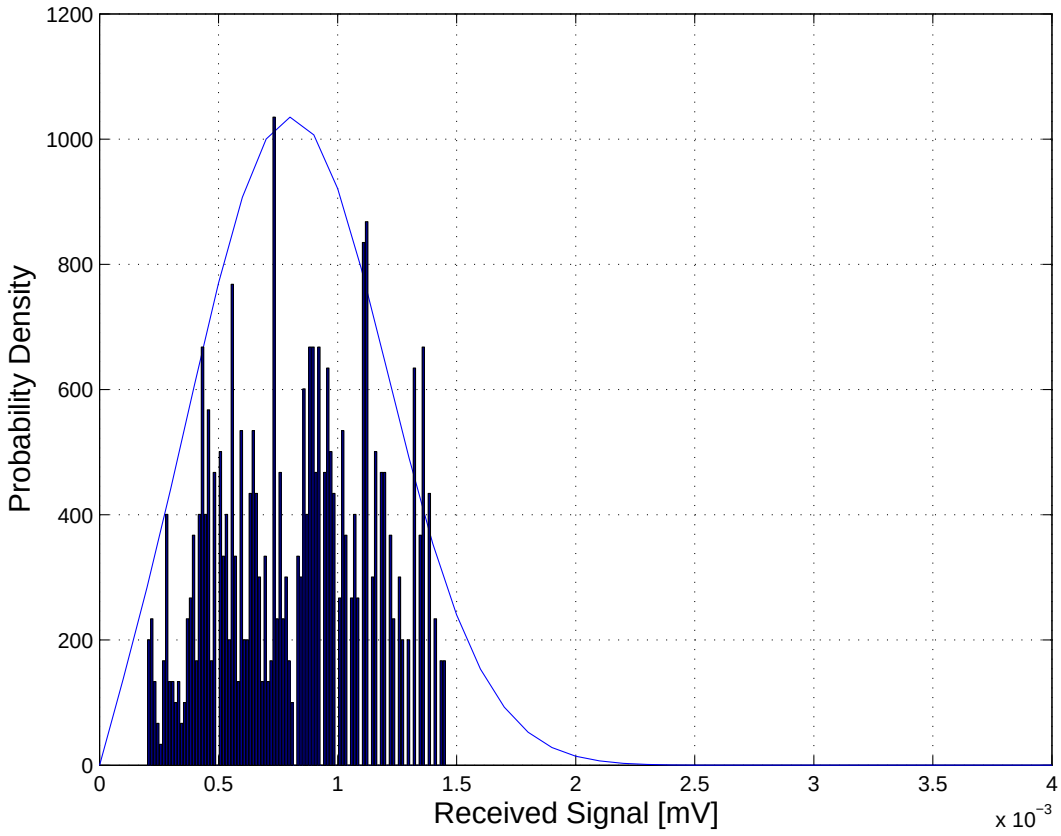

Figure 2.14. Match between experimental results and a Ricean pdf with K=1.43.

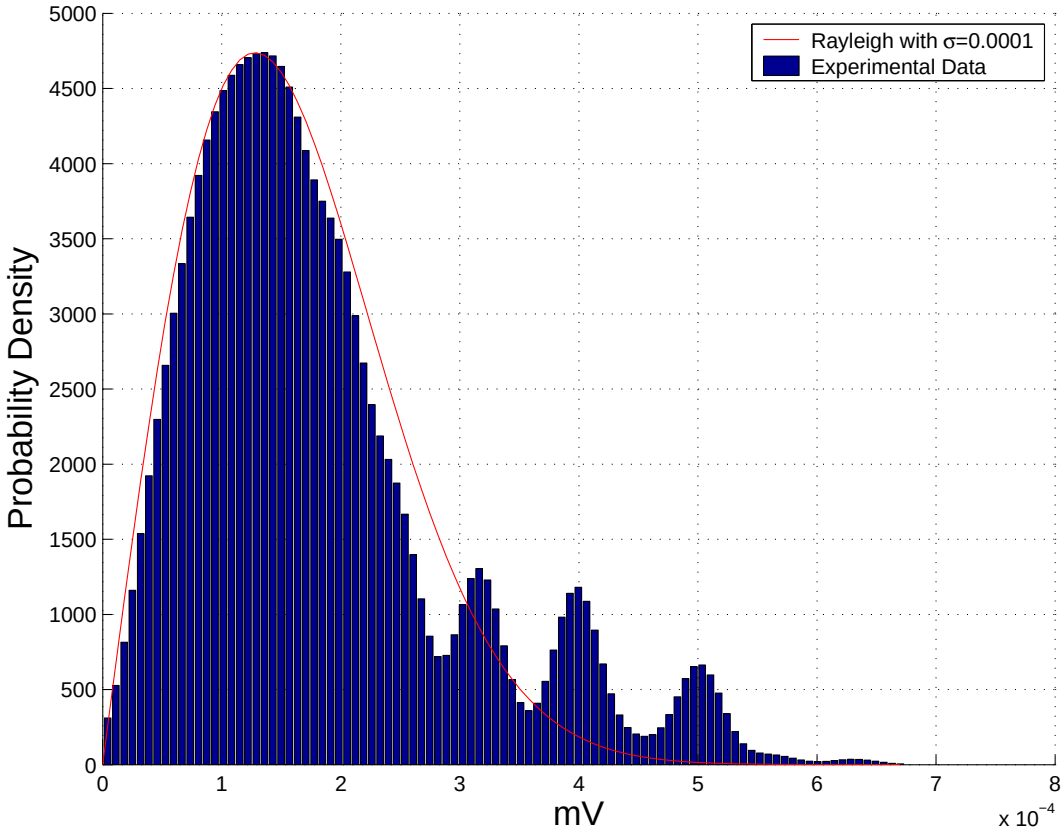

Figure 2.15. Experimental verification of Rayleigh Fading with Telos hardware.

maintained in both cases at all times. Figure 2.14 shows a reasonable match given by our experimental results. It should be noted that Figure 2.14 refers to an experimental setting corresponding to Figure 2.13, where both motes are on the turntable so that the dominant component is fully static, in harmony with the assumptions behind the Ricean model.

From an analytical standpoint, the Rayleigh model provides a wieldy tool for characterizing fading in a probabilistic fashion. The same cannot be said of the Ricean model, which provides a probability density featuring Bessel functions. The key to Rayleigh's simplicity is obviously the basic assumption of the absence of

a direct line-of-sight path. Figure 2.15 shows the probability density measured during a turntable test with Telos motes. In this case, one mote was placed on the turntable, and the other was hidden inside a metal file cabinet. On average, they were 6m apart. It should be noted that there remains a trace of the direct path, as can be observed by looking at the bumps in the tail of the experimental data.

2.2.5 A Probabilistic Link Model

Packet transmission is typically modeled as a probabilistic process. We now look at an interesting link model suggested in [30, 63], we define a collection of k Bernoulli random variables X_i that take on 1 if the i-th packet is received correctly and 0 otherwise, are independent and identically distributed. The probability of correct packet reception p_r may be written as $\mathbb{E}(X_i)$. If we perform k packet transmissions, the average number of correctly received packets is $\frac{1}{k}\sum_{i=1}^{k} X_i$; by the weak law of large numbers, this can be approximated by $\mathbb{E}(X_i)$ as $k \to \infty$.

Our focus is on Mica2 hardware, which is the basis for the hardware implementations described in chapter 3. The radio used by this platform, CC1000, performs non-coherent BFSK modulation and uses NRZ (No Return to Zero) encoding. If p_e is the bit error probability, we can express the probability of successful reception of a packet in our system as

$$p_r = (1 - p_e)^{8l}(1 - p_e)^{8(s-l)} = (1 - p_e)^{8s} \tag{2.13}$$

where l and s respectively indicate the length of the preamble and the frame size (preamble, payload, and CRC) in bytes. This model assumes independent fading for each bit, which makes sense in the presence of dynamic fading. Despite this, we will show how this model gives reasonable results in our static context as well.

In a dynamic Rayleigh fading channel, we can write the bit error probability as

a function of the bit energy to noise ratio $\frac{E_b}{N_0}$ as

$$p_e = \frac{1}{2}\exp\left(\frac{E_b}{2N_0}\right) \tag{2.14}$$

Thus, we can write p_r as

$$p_r = \left(1 - \frac{1}{2}\exp\left(\frac{E_b}{2N_0}\right)\right)^{8s} \tag{2.15}$$

CC1000, as well as most radios, provides a measure of the RSSI, which can be obtained from the SNR γ if we know the noise floor. Empirical measurements show that the noise floor for Mica2 is between -100dBm and -105dBm. From communication theory, the noise floor can be computed analytically as

$$P_n = (F+1)kT_0B_N \tag{2.16}$$

where F is the noise figure, k is the Boltzmann's constant, T_0 is the ambient temperature, and B_N is the noise bandwidth. CC1000 has a noise bandwidth of 30kHz and a noise figure of 13dB, which is only indicative of the radio chip itself; we should expect a higher value for the whole device. At a temperature of 300K the theoretical noise floor (using the radio noise figure) is -115dBm. This is a good match with our empirical value if we consider a higher noise figure and the ±6dB accuracy on the RSSI readings.

The SNR γ is related to $\frac{E_b}{N_0}$ by

$$\gamma = \frac{E_b}{N_0}\frac{R}{B_N} \tag{2.17}$$

where R is the bit rate, which is equal to 57.6kb/s for Mica2. The probability of correct reception can therefore be expressed as

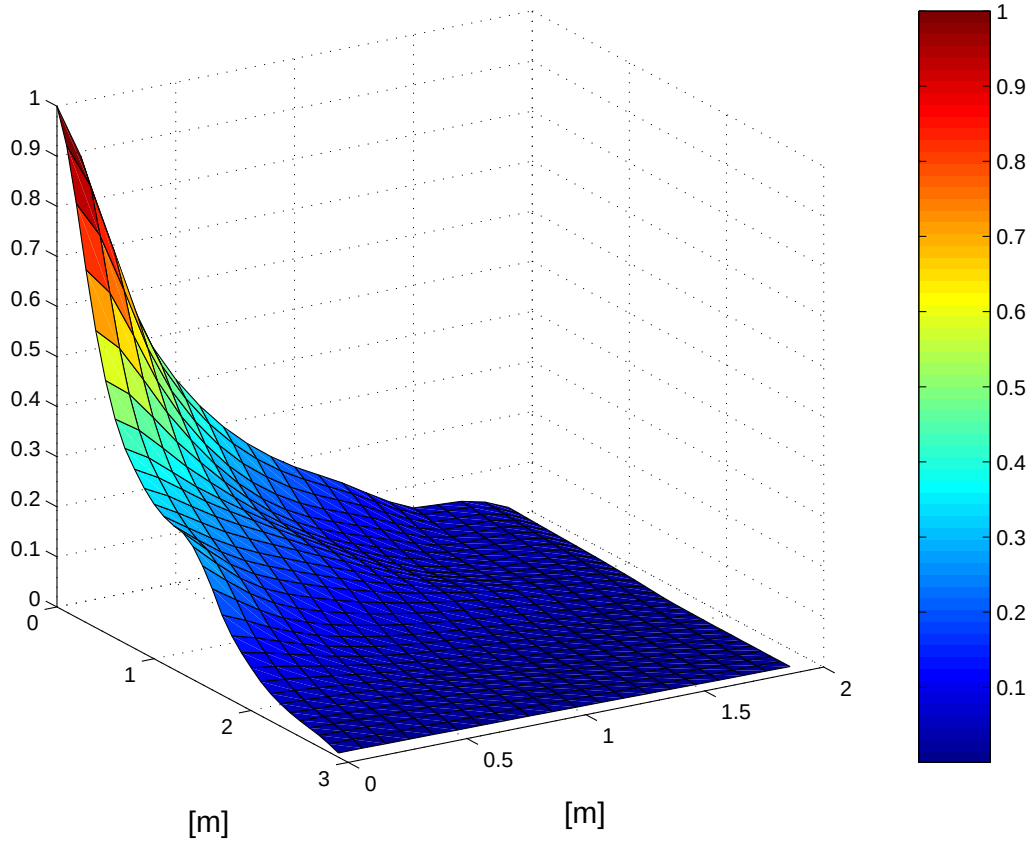

Figure 2.16. The spatial autocorrelation function pertaining to Figure 2.3.

$$p_r = \left(1 - \frac{1}{2}\exp\left(-\frac{\gamma}{3.84}\right)\right)^{8s} \tag{2.18}$$

2.2.6 Spatial Correlation of the Fading function

It is interesting to look into the properties of the two-dimensional autocorrelation function of the fading function. Starting from the output of our simulator, we compute the spatial autocorrelation of the fading function as

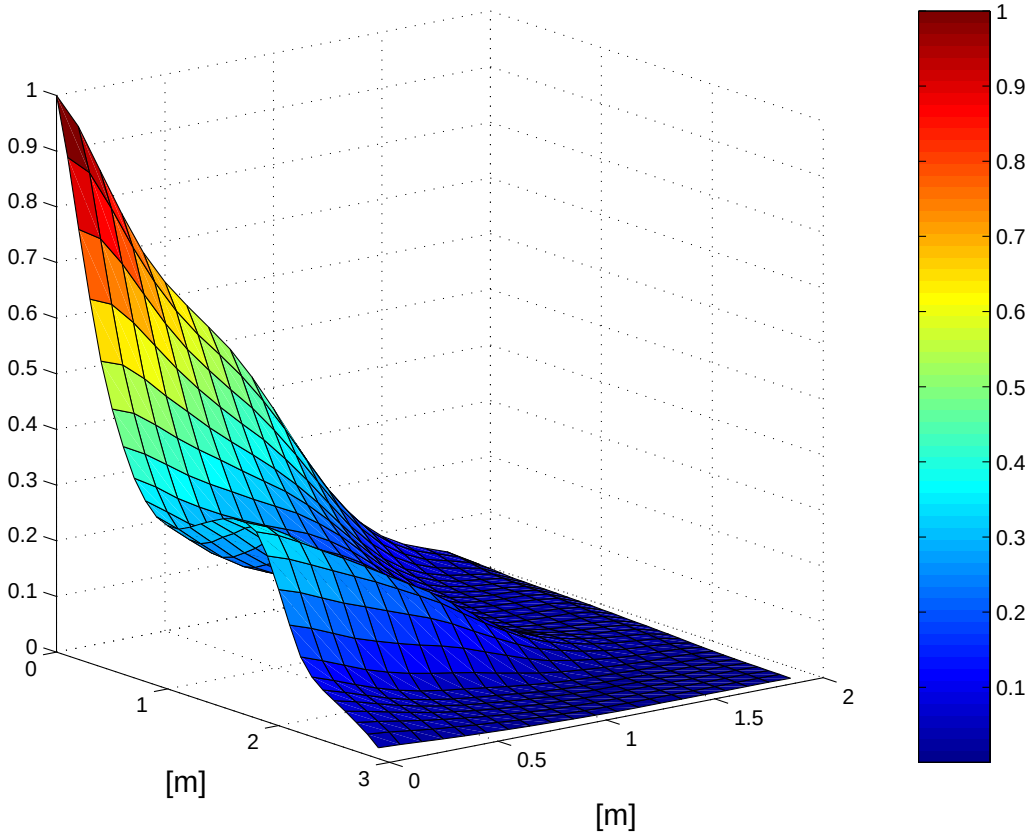

Figure 2.17. The spatial autocorrelation function pertaining to Figure 2.4.

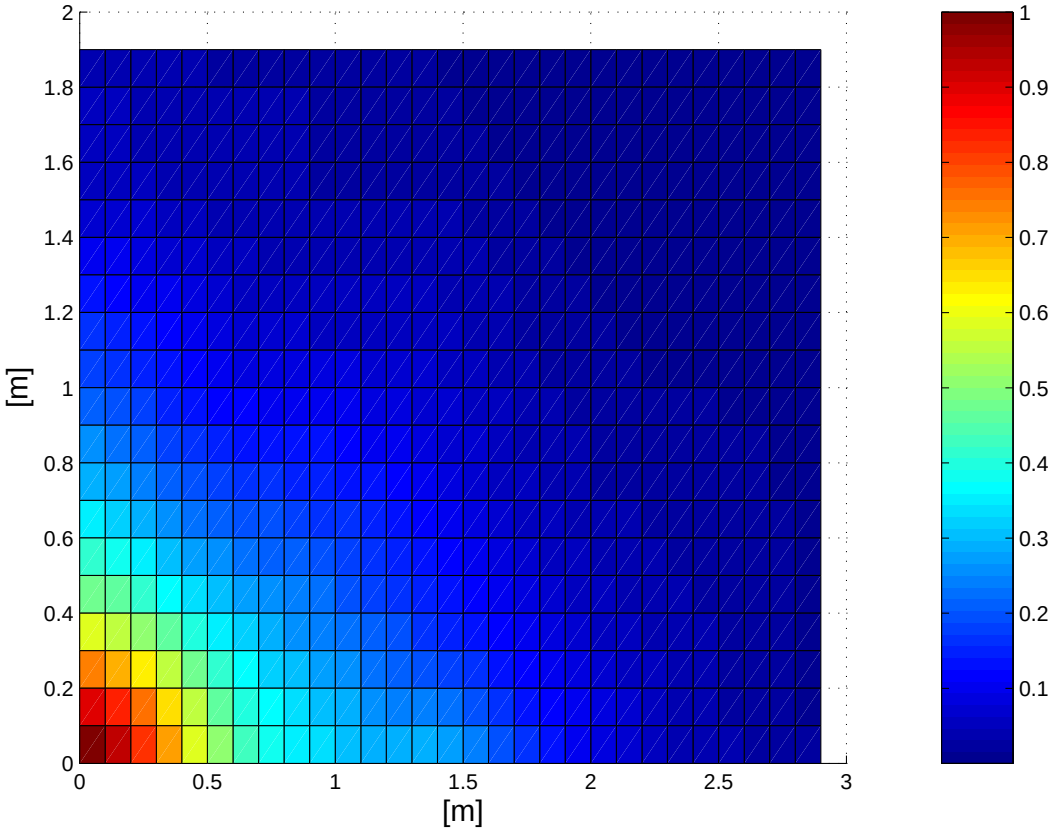

Figure 2.18. Planar projection of Figure 2.16: each cell measures 10×10cm, and $\frac{\lambda}{2}$ is 34.5cm.

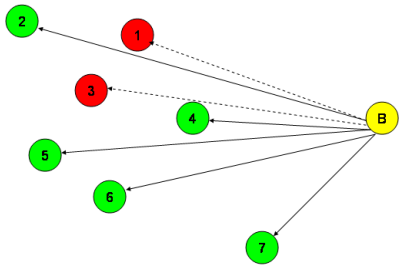

Figure 2.19. A typical sensor network topology.

$$R(n_1, n_2) = \sum_{k_1=0}^{N_x-1} \sum_{k_2=0}^{N_y-1} f(k_1, k_2) f(k_1 - n_1, k_2 - n_2) \tag{2.19}$$

The normalized spatial autocorrelation function for the fading function in Figure 2.16 (resulting from a $3m \times 2m$ topology with the transmitter placed at (1.5m, 0.1m)) shows how the signal strength is highly correlated over spatial variations of less than $\frac{\lambda}{2}$. It is also worth observing that the spatial autocorrelation exhibits a decreasing behavior, which indicates that in this case having terminals $\frac{\lambda}{2}$ apart is not the same as having them $N\frac{\lambda}{2}$ apart. A similar behavior is displayed by Figure 2.17, pertaining to a $3m \times 2m$ topology with the transmitter placed at (1.5m, 0.1m). This confirms that MIMO systems where antennas are spaced by $\frac{\lambda}{2}$ undersample the fading function; the addition of mobility can definitely help sensor nodes exploit the properties of the spatial autocorrelation of the fading function and maximize the spatial diversity benefit.

2.3 Impact on Sensor Network Design

Our approach based on the fading function has interesting applications in the context of sensor network design. Let us consider the situation depicted in Figure 2.19, which represents an indoor deployment of sensor nodes. The base station (B) needs to obtain sensor readings from the sensing nodes, and sequentially polls them for data; the natural structure of this network follows therefore a star topology. Owing to the layout of the environment, the links between the base station and nodes 2, 4, 5, 6, and 7 are mapped to good channel gains by the fading function, whereas the links to nodes 1 and 3 are mapped to poor channel gains. We are basically reducing the mapping given by the fading function to a binary mapping where nodes can either be in a good or a bad fading spot by introducing a threshold. This simple example is indeed a good model for many practical deployments. We assume the sensing nodes to be severely resource-constrained, primarily in terms of energy, but also in terms of storage capacity as well as computing power. However, we also assume the base station to be less resource-constrained than the nodes. Typically, a base station is bridged to a gateway with superior computing resources. In this scenario, the goal is clearly the minimization of the energy expenditure at the sensing nodes. Sending packets is costly, so the sending nodes should send as few packets as possible, and most importantly packet loss between the nodes and the base station should be carefully avoided. Processing is also a delicate issue: the sensing nodes should be required to perform lightweight processing tasks and only if strictly necessary. Ideally, they should sleep as much as possible, and only wake up to receive the command from the base station, obtain the required sensor data from their ADCs or digital sensors, send the data, and go back to sleep. It should be noted that receiving is almost as costly as sending, but low power listening techniques can be used [51] along with a heavy exploitation of sleep modi. In this

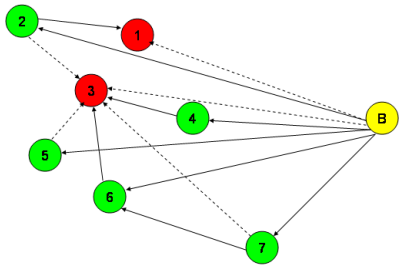

Figure 2.20. The fading function maps each link to a different channel gain: nodes can be in good fading spots with respect to other nodes, even though they are in a bad fading spot with respect to the base station.

scenario, we focus on minimizing the probability of packet loss from the sensing nodes to the base station due to noise. By noise we indicate the effects of multipath fading, as well as other effects such as shadowing. Since polling is used, we can assume the absence of interference.

In the setting shown in Figure 2.19, when the base station polls nodes 1 and 3, packet loss is very likely to occur. The base station is not as energy-constrained as the sensing nodes; if it does not receive the requested information, it may retransmit the request. However, if the topology of the room causes nodes 1 and 3 to be in a particularly deep fade, retransmissions may be useless. Multihop routing is usually proposed as the natural solution for problems of this nature. Since the fading function maps each link to a different channel gain, we might have a situation like the one depicted in Figure 2.20, where the base station can reach 1 and 3 leveraging on the other nodes. The literature presents countless routing strategies, and lightweight solutions exist for the implementation of multihop routing schemes

in sensor networks. This approach, however, would not allow us to meet the overall goal of minimizing the energy expenditure at the sensing nodes. The nodes acting as relays would have to give up on their sleep to enforce the chosen routing scheme, and this would be costly in terms of processing and communication. Nodes regularly acting as routers would use up their battery power much sooner than they would if they were not overburdened by having to carry out extra work. Ironically, in this scenario nodes sitting in a good fading spot with respect to the base station would die sooner than nodes located in bad fading spots.

It would be desirable for the individual sensing nodes to be reached by the base station in a single hop, but some nodes can be beyond the reach of the base station, and for the ones within the radio range the mapping given by the fading function can dictate prohibitive channel conditions as long as everything is static. The introduction of mobility in the network could definitely help, but the question is where it should be introduced. Making the sensing nodes mobile is out of the question, as their limited resources may not be wasted for motion. Making the base station mobile introduces other problems: typically, a base station is wired to a gateway bridging the sensor network to an external infrastructure. In the case of Mica2 or MicaZ motes, typical base stations are serial programming boards connected to a computer (MIB510), Ethernet programming boards (MIB600) or more sophisticated devices such as Stargate. In the case of the Telos platform, any mote can be a gateway, so long as it is directly connected to the USB port of a host machine. With virtually all the currently available platforms, mobility in the base station would necessarily imply a mobile gateway, which is a costly solution, and is only viable with gateways sophisticated enough to be able to establish a wireless link with a remote host. Our solution is adding mobility to an agent operating as an additional interface between the base station and the sensing nodes: a *mobile*

sink with not-so-strict energy-constraints moving along a predetermined trajectory and polling the sensing nodes.

In Chapter 3, we will detail this concept and show how beneficial it can be in actual sensor network deployments.

CHAPTER 3

A CROSS-LAYER APPROACH TO MITIGATE MULTIPATH FADING

3.1 The Mobile Sink Concept

The introduction of a mobile sink in a sensor network like the one shown in Figure 2.19 is a simple, powerful tool to enhance the signal-to-noise ratio. The mobile sink is a node with mild energy constraints (comparable to the actual base station) which is capable of moving along a trajectory which can be arbitrary, predetermined, or learned on an ongoing basis. During its motion, it may also be in bad fading spots with respect to the base station. However, its mobility guarantees the possibility of establishing an at least intermittent link to the base station, and we assume the possibility of temporary data storage in a non-volatile memory and bulk transfers whenever good connectivity to the base station exists: its mild resource constraints make it possible for the sink to guarantee data delivery to the base station. The advantage of having a mobile agent within a network is clear if we think in terms of fading function mappings. In Figure 3.1, nodes 1 and 3 are in a bad fading spot with respect to the base station (B). However, one of them happens to be in a good fading spot with respect to the mobile sink (S), which can easily reach it. The mobile sink then finds itself in the position indicated in Figure 3.2, and is now able to reach both 1 and 3. Of course, the list of nodes that the mobile sink can reach changes as the position of the mobile sink changes. This simple idea exploits spatial diversity in a dynamic and continuous fashion, contrarily to multi-antenna

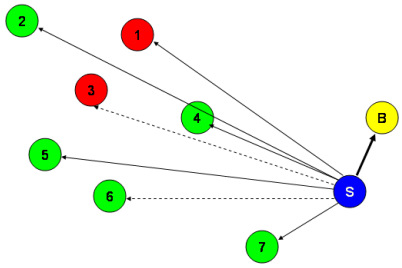

Figure 3.1. The solid arrows represent good links, whereas the dashed ones indicate bad links (the receiver is in a bad fading spot with respect to the transmitter) Nodes 1 and 3 are in a bad fading spot with respect to the base station, but 1 is in a good spot with respect to the mobile sink.

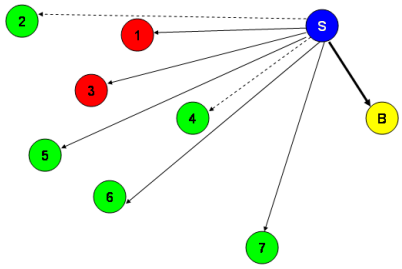

Figure 3.2. Now the mobile sink has changed its position and both 1 and 3 are in good fading spots with respect to it. Of course, nodes with a good link with the base station may have a bad link to the mobile sink depending on its position (which is the case with node 6 here).

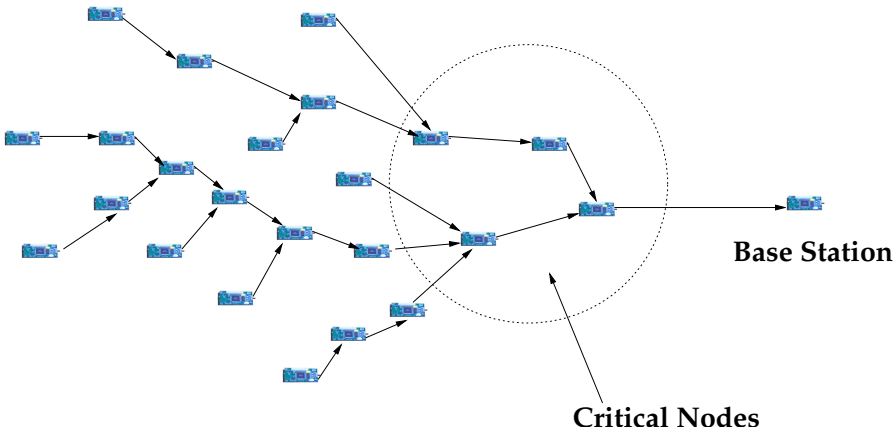

Figure 3.3. A uniform energy consumption pattern should avoid the depletion of the resources of nodes located in the vicinities of the base station.

system which generally undersample the fading function. Placing multiple antennas on a mote is unpractical, as the antennas would have to be spaced by at least $\frac{\lambda}{2}$. Another MIMO solution could be having multiple base stations in a Master/Slave(s) architecture where the Master base station is wired to the gateway and is in charge of fusing the data coming from the Slave base station(s). The advantage of our scheme is the dynamic element: a mobile agent can exploit the topology-dependent irregularities of the fading function and the properties of its spatial autocorrelation.

The problem of efficiently funneling sensor data from the nodes to the base station in the sensor network is usually seen as a routing problem. Oversimplifications in the analysis of the problem lead many authors to support techniques based on multihop routing over many short-hops. Routing over fewer (possibly longer) hops carries indeed many advantages [25]; from our point of view, the most interesting are the reduction of energy consumption at the sensing nodes, the achievement of a better energy balancing, the more aggressive exploitation of sleep modi, and the lack of route maintenance overhead. If the sensing nodes only occasionally need to act as relays, they can sleep longer and only consume energy to make their own

data available. Also, we can avoid scenarios where a few nodes in a critical area in the proximity of the base station (Figure 3.3) have to relay all the traffic and use up their battery power prematurely. Last but not least, the computational burden associated with keeping track of the routes can be eliminated. Having as few hops as possible is desirable; with the addition of a mobile element bridging the base station to the sensing nodes, data retrieval can be done in a single hop between the sensing nodes and a mobile sink, and the sensing nodes may sleep any time they are not called upon to send their data to the sink. This observation paves the way to the introduction of a simple data collection and routing algorithm, MobiHop, which we have designed and successfully implemented with Mica2 hardware.

Motion control and planning mainly depend on the application, the number of sensing nodes, and the deployment area. For many applications, it is an unnecessary complication, and arbitrary motion works reasonably well.

The natural question at this point is how we can determine the quality of a link. Popular techniques include the use of RSSI and LQI (Link Quality Indicator), and passive estimation (snooping) [59]. We believe that a metric should be chosen with special regard to the minimization of processing at the sensing nodes.

3.1.1 Signal Strength as a Link Quality Estimator

Received signal strength is a good tool, but its unreliability as a link quality estimator is often underscored. In the presence of interference, it is indeed unreliable: the received signal strength may be very good, but collisions may still cause a heavy packet loss.

Recent developments in the sensor network community call for the adoption of CC2420-based platforms such as MicaZ from CrossBow or Telos from MoteIV. This new radio from Chipcon implements IEEE 802.15.4 (ZigBee) and offers an

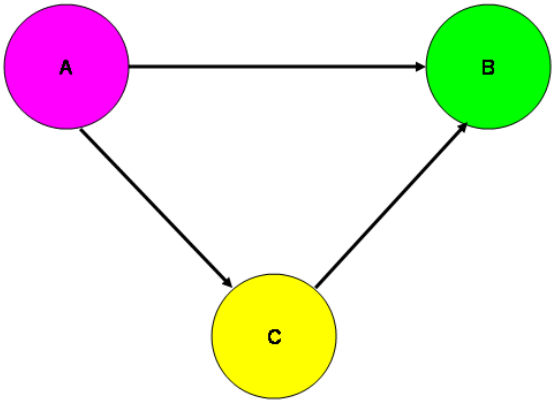

Figure 3.4. A minimal sensor network with ongoing communication. We assume the presence of a MAC scheme to avoid collisions.

estimate of the BER called Link Quality Indication (LQI). It has been shown that LQI exhibits a very good correlation with packet loss, and is therefore a much better link quality indicator (hence its name) [41]. However, one of the contributions of the present work is to show that RSSI is a reasonable metric if it is processed correctly, and if interference can be distinguished from noise. Given that LQI is a superior metric, it should not be forgotten it is only made available by 802.15.4-compliant devices. It therefore makes sense to make the most out of RSSI.

In a setting where interference is not an issue, proper processing can make RSSI a reliable estimator of link quality. In the simple wireless network shown in Figure 3.4, node C is receiving packets from node A, and node B is receiving data from both A and C. In the case of node C, things are simple: all the signal strength samples can be processed together, as they pertain to the same link. However, it would be meaningless for B to process signal strength samples pertaining to packets from A along with ones related to packets from B, as the two links A $\rightarrow$ B and C $\rightarrow$ B are mapped to different channel gains by the fading function. As long as separate signal strength buffers are maintained for each active link, the reliability of the RSSI is reasonable.

Things can be improved by fusing RSSI monitoring and packet snooping. Thanks to the broadcast advantage of radio transmissions, the sensing nodes receive all the requests made by the mobile sink. Using low-power listening techniques, a sensing node can pick them up and measure the associated signal strength, which corresponds to the channel gain to which the fading function maps the link between the sink and the sensing node. Moreover, if the packet has a sequence number, the sensing node can do some simple math to keep track of the packet loss rate associated with each link, and use it to correct the link quality estimate given by RSSI monitoring. This approach also helps the sensing nodes make sense out of

those case in which interference provides the illusion of a high signal strength level.

We have implemented a number of applications in TinyOS where RSSI monitoring and passive snooping are applied. TinyOS facilitates the applicability of these idea by providing Active Messaging. Virtual links between motes can be established by the programmer using a concept that is similar to TCP ports from the point of view of the application programmer, but is implemented at a fairly low-level with a special byte in the packet header. An interesting application that exploits the mobile sink architecture is MobiHop, a single-hop data retrieval application for sensor networks.

3.2 MobiHop: an Energy-Aware Protocol for Data Collection

MobiHop is a simple, lightweight data retrieval protocol for sensor networks based on the presence of a mobile sink. The main idea is to exploit opportunistic transmissions induced by spatial diversity to route the data from the sensing nodes to the mobile agent in a single hop. The mobile sink polls the sensing nodes to obtain their sensor readings on behalf of the base station. The sensing nodes operate in a low-power listening mode. If a node receives a request from the mobile sink meant for a different node, the request is not processed and the RSSI information from the radio is sampled to assess the channel. Each node keeps an overall running average of the RSSI information obtained by snooping on requests addressed to the other nodes, and a running average of the latest N RSSI samples (buffer average), which provides a filtered look at the current channel conditions.

Transmission decisions can be performed in various ways, such as comparing the latest RSSI sample to the overall running average, comparing the buffer average to the overall running average, or simply comparing the latest RSSI sample or the buffer average to a fixed threshold. Further, the buffer and overall average can be

used to dynamically adapt the transmission power: the sensing nodes can transmit opportunistically and exploit the best channel conditions. The mobile sink can then relay the data to a base station if possible, or store the data and relay it at a later time with a bulk transfer.

The first question that comes to mind is how to manage the motion of the mobile sink. Various options exist, and their validity depends on the application. In principle, the best solution is having an intelligent vehicle controlled by the sink itself. The sink can start at an arbitrary position, poll the sensing nodes, and assess the effectiveness of the position by the number of nodes that could be reached and decided to respond. It can then move to different positions until a position or an interval are found where all nodes can be reached. We did assume the mobile sink not to be as resource-constrained as the sensing nodes; however, some applications may rule out motion control as too expensive, impractical, or even unnecessary. Indeed, simpler solutions exist. A predetermined trajectory is a good option if it can be obtained without using the resources of the mobile sink. For all practical purposes, an arbitrary trajectory provides an excellent compromise between performance and cost/simplicity.

The main goal of this scheme is minimizing or even avoiding packet losses between the sensing nodes and the mobile sink. However, we are also interested in hearing from all nodes. For small numbers of sensing nodes the polling process may be implemented with a simple round robin, but as the number of sensing nodes increases more sophisticated scheduling schemes can be used. A good tradeoff between simplicity and effectiveness can be reached by a weighted round robin where the sensing nodes have an adaptive weight which is decreased upon their response.

3.2.1 Probabilistic Considerations on MobiHop

If the mobile sink polls the sensing nodes with a simple round robin, each node measures RSSI values which depend on its position with respect to the mobile sink. These measurements are determined deterministically by the topology of the environment. If we abstract ourselves from a particular topology, the probabilistic viewpoint tells us that the signal amplitude follows a Ricean distribution in the presence of a line-of-sight and a Rayleigh distribution in the absence thereof. Let us suppose that the nodes transmit if the measured RSSI value associated with the request exceeds the running average of the channel conditions. In probabilistic terms, they respond if the signal level associated to the measured RSSI exceeds the mean of the distribution. Since the Rayleigh fading model pertains to a worst-case scenario (no line-of-sight) and is more analytically tractable than the Ricean model, we will assume that the field values corresponding to the RSSI samples measured by the sensing nodes follow a Rayleigh distribution with variance σ^2. Let R^2 be a random variable modeling the RSSI values measured by a node; the mean of the related signal amplitude is given by

$$\mathbb{E}[R] = \sigma\sqrt{\frac{\pi}{2}} \tag{3.1}$$

The complementary cumulative distribution function for a Rayleigh-distributed random variable is given by

$$\mathbb{P}(R \geq r) = \int_0^R p(r)dr = \exp\left(-\frac{r^2}{2\sigma^2}\right) \tag{3.2}$$

So the probability of a node responding to a request is

$$\mathbb{P}\left(R \geq \sigma\sqrt{\frac{\pi}{2}}\right) = \exp\left(\frac{\pi}{4}\right) \approx 0.456. \tag{3.3}$$

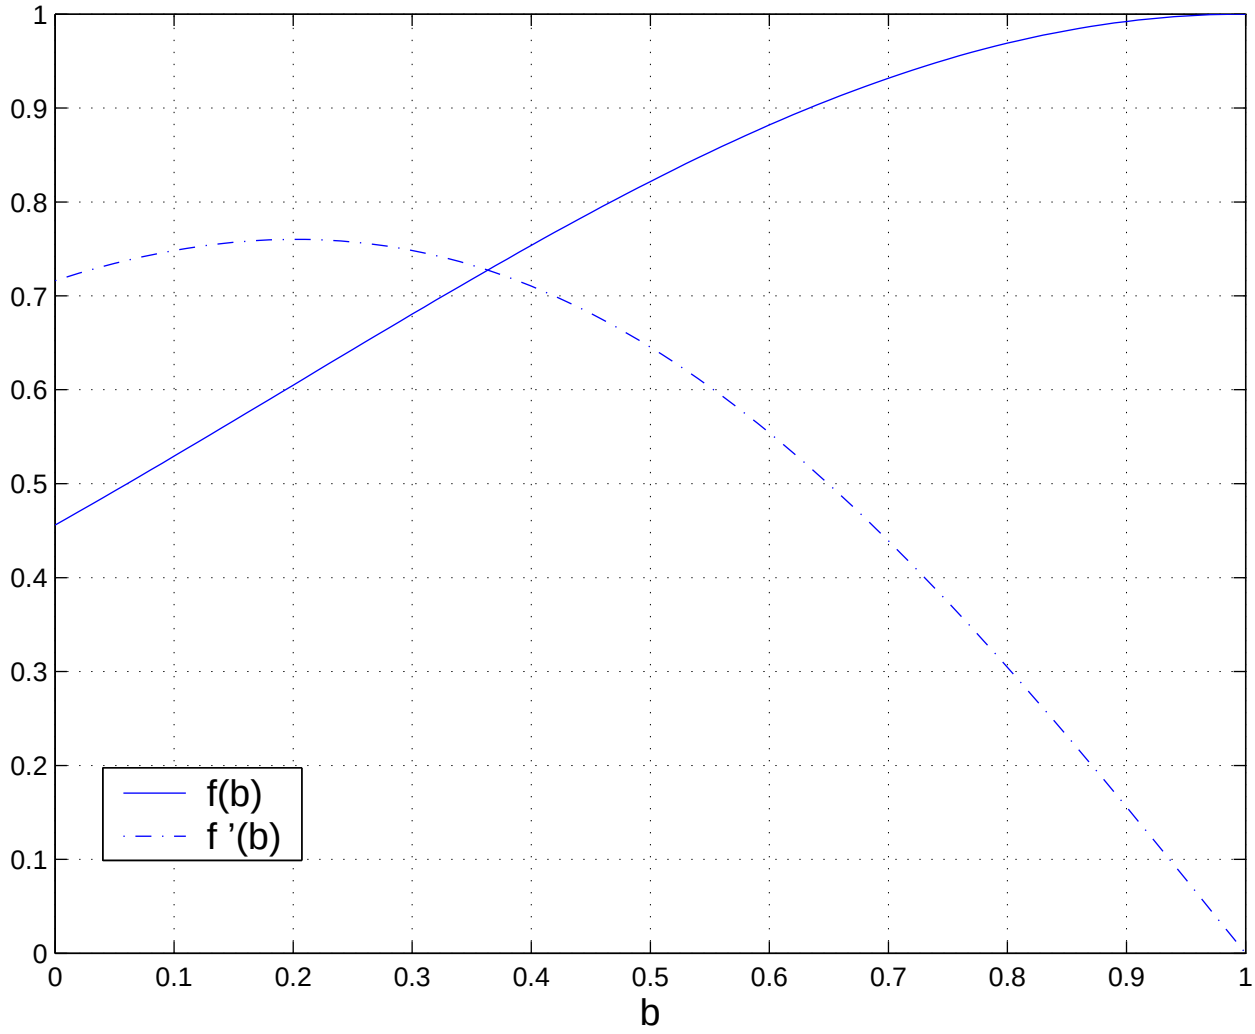

Figure 3.5. A small correction factor for the threshold can significantly change the probability of response.

With the Rayleigh model, the sensing nodes reply about 46% of the times independently of the variance of the distribution, *i.e.* independently of their position with respect to the sink and the motion of the latter. Ricean distributions have stronger symmetry properties than Ricean distribution, so it makes sense to assume that on average the sensing nodes are expected to reply 50% of the times. We define the efficiency of MobiHop as the ratio of responses to requests; efficiency is $\frac{1}{2}$ in this case, and can be increased by lowering the threshold with respect to the running average. If we lower the threshold by a factor $b \in [0..1]$ of the mean of the Rayleigh distribution, we get

$$\mathbb{P}\left(r \geq \sigma\sqrt{\frac{\pi}{2}}(1-b)\right) = \exp\left(-\frac{\pi(1-b)^2}{4}\right) = f(b) \tag{3.4}$$

Figure 3.5 shows that conservatively small threshold lowering factors can nonetheless improve the efficiency in terms of ratio of number of responses to number of requests due to the fairly good sensitivity of the function $f(b)$ with respect to b.

However, comparisons to a fixed thresholds are also an option, with the advantage that the processing overhead associated with the computation of the running averages of the channel conditions becomes unnecessary for transmission decisions and only remains important for transmission power regulation (and can be avoided altogether on nodes with extremely low computing power).

Figure 3.6 compares the analytical results from 2.18 with our experience in indoor environments. The analytical curve is indeed a bit pessimistic, but this is because of the underlying Rayleigh fading channel model, which assumes the absence of a direct path. The takeaway is that -80dBm (25dB above the noise floor) is a conservative calibration for the threshold of correct reception (it should not be forgotten that there is a ±6dB accuracy in the RSSI measurements).

From [23], in a Rayleigh Fading network, the probability of successful reception

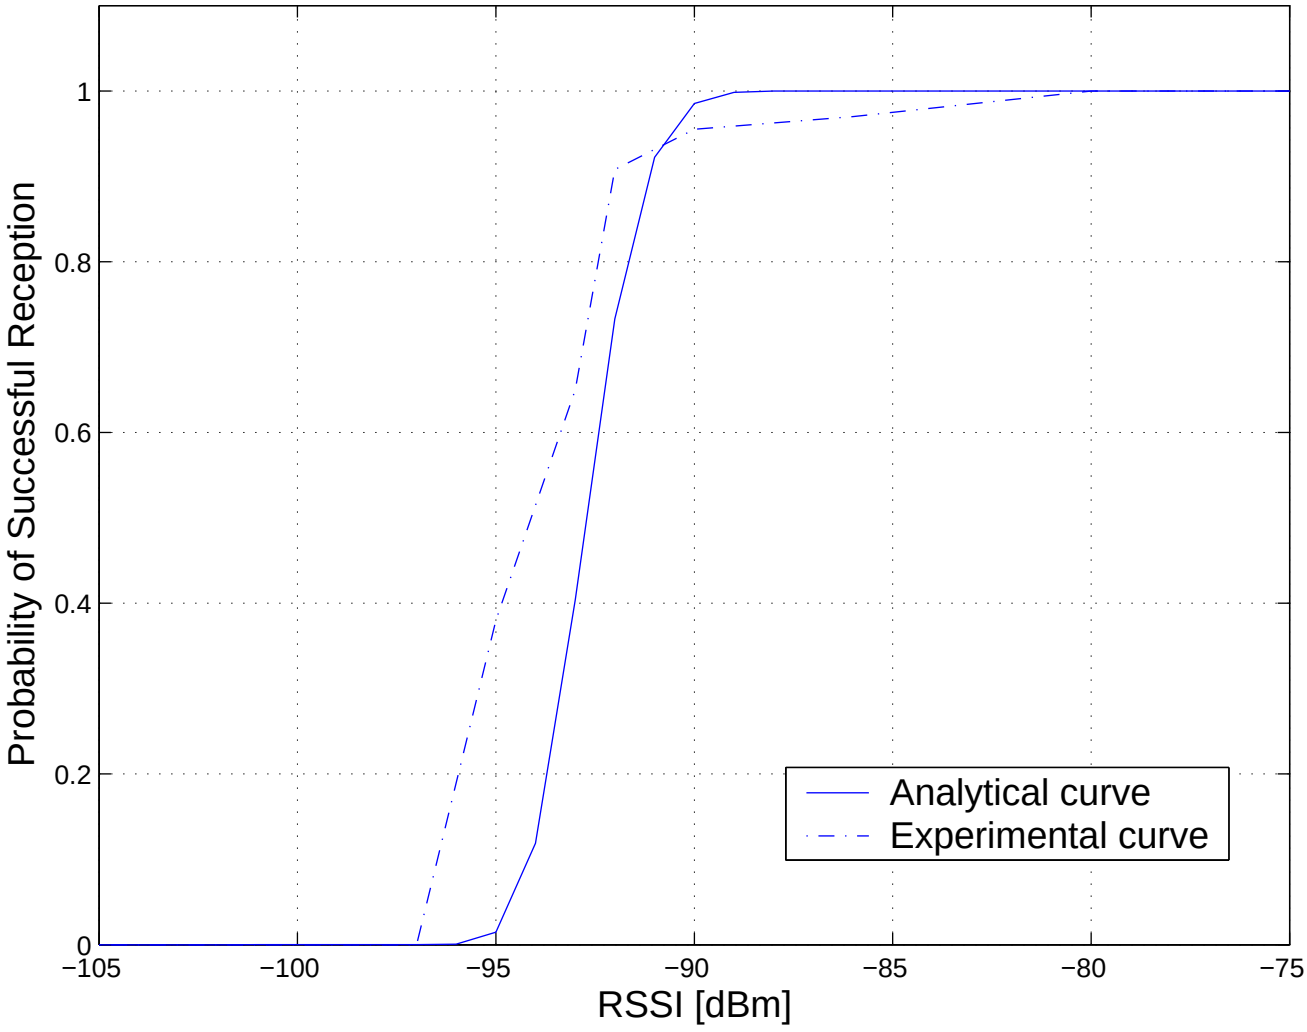

Figure 3.6. Comparison of analytical and experimental results for the probability of successful packet reception. The experimental curve corresponds to a transmission power of 0dBm, and has been obtained by averaging the received power at given distances from a fixed transmitter.

in the absence of interference is expressed as the probability of the SNR exceeding a threshold θ:

$$p_r = \mathbb{P}[\gamma > \theta] = \exp\left(-\frac{\theta N_0}{\overline{P_{rec}}}\right) \tag{3.5}$$

where $\overline{P_{rec}}$ indicates the mean of the received power, which in the Rayleigh model is exponentially distributed, and N_0 denotes the noise power. We can use this model to predict the efficiency of MobiHop given a particular threshold. If we choose -80dBm according to the arguments above, the threshold θ corresponds to 25dB. Figure 3.7 shows that the choice of θ=10dB, often made in the literature, corresponds to an analytical curve which compares favorably with our Mica2 experimental results. The same Figure also shows the curve for θ=25dB, which is the margin which we have empirically chosen to predict a correct reception. This curve looks at the probability of the SNR being more than 25dB above the noise floor for a particular mean received power. We can use our turntable as a tool to obtain a fading channel and gain some information on the expected efficiency of MobiHop. For instance, if a sensing node is placed in a spot where the mean RSSI is -70dBm, the model predicts that the received signal will be 25dB above the noise floor 90% of the time: in a Rayleigh fading channel, we expect an efficiency of 9 responses every 10 requests. Hence, a natural way to assess the efficacy of our protocol is a hardware implementation where the mobile sink is placed on a turntable.

3.3 A hardware implementation of MobiHop with Mica2

We have implemented MobiHop with Mica2 hardware with 5 sensing nodes, a mobile sink, and a base station. A separate virtual channel is defined for each of the nodes; the virtual link between the mobile sink and node a is identified by an Active Message type which equals the address of node a. Each sensing node can recognize

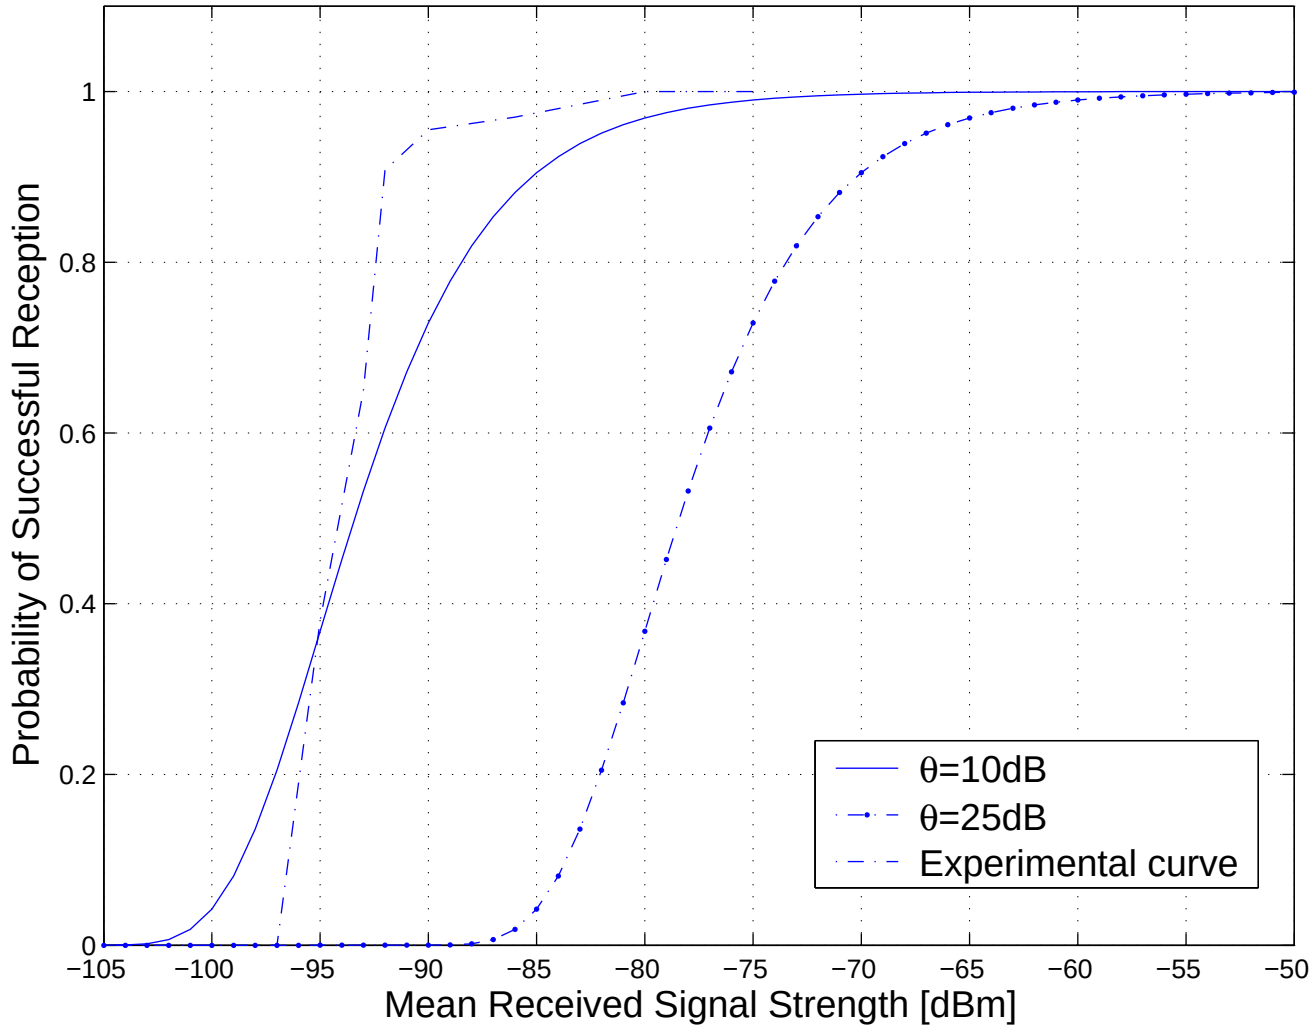

Figure 3.7. An assessment of the efficiency of MobiHop in terms of responses to requests ratio obtained from a probabilistic model.

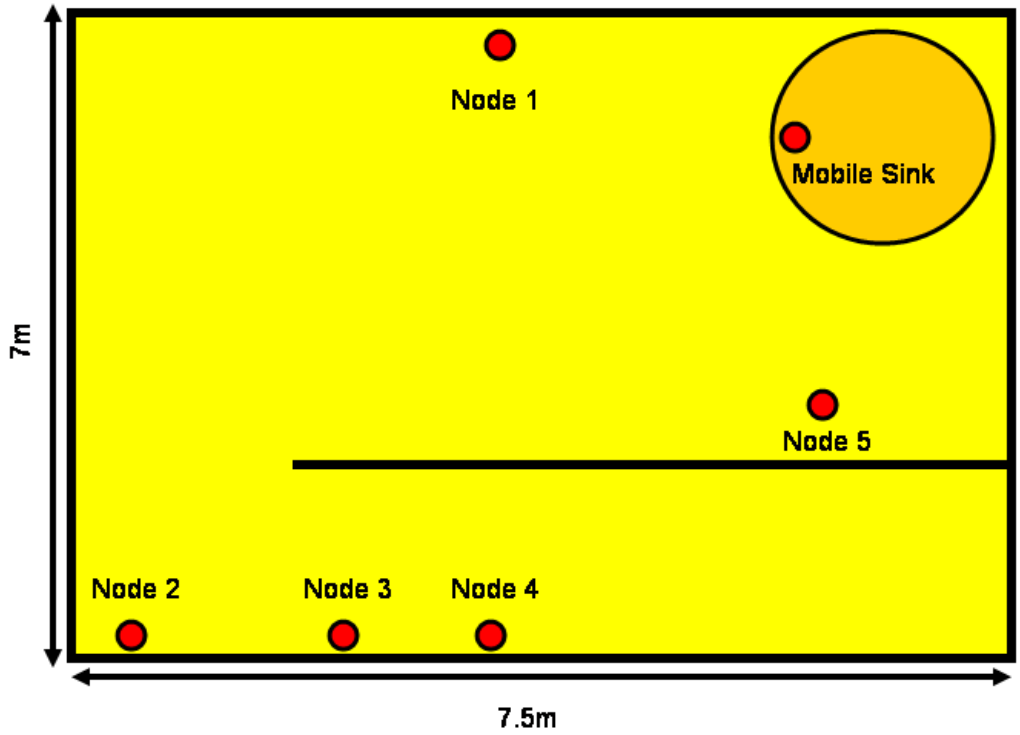

Figure 3.8. The layout of the indoor environment used for the experiment

and react to all 5 types, but only processes packets whose type corresponds to its address, while using the others to sample the RSSI from the CC1000 and feed it into the running average. It should be noted that, independently of the destination for which a packet is intended, any node that receives it can use it to assess its link to the sink in that moment, and use it for the computation of the running average. The 5 nodes are deployed in an indoor environment along with the mobile sink. To check the efficiency of MobiHop in a fading channel and compare it with the analytical results for Rayleigh fading channels, we place the mobile sink on our motorized turntable. The layout of the room is shown in Figure 3.8. The base station is placed in the vicinity of the turntable. This setting is representative of practical situations that might arise from the deployment of sensor nodes in indoor environments. Building monitoring (Smart Buildings) is an example of an application that could exploit MobiHop.

We have seen that having the sensing nodes decide whether to transmit on the sole basis of a comparison of the current RSSI against a running average is inefficient, and leads to an expected throughput of 0.5 responses per request. If the mean of the field distribution seen by a node is relatively high (*i.e.*, it corresponds to RSSI values in the range [-60, -50]dBm), things are even worse because of the typically small variance of the distribution: the node would refrain from transmitting even in the presence of a perfectly good channel. We showed that the use of a correcting factor can improve efficiency. In practice, our hardware provides us with an affine mapping of the RSSI in dBm; a possible approach is applying a threshold directly to this quantity, in order to minimize the processing at the sensing nodes. However, a fixed threshold is the simplest option. In our implementation, the sensing nodes use a fixed threshold and keep a moving average of the channel conditions to implement an opportunistic power saving scheme: the transmission power is dynamically adapted

to the channel conditions.

In Mica2, the affine mapping between the CC1000 RSSI register values and dBm values depends on the battery voltage. By using a fixed register value yielding -80dBm with new batteries (3.2V), the threshold will be about 4dB lower once the batteries provide 2.8V. It is of course possible to continuously adjust the threshold according to the battery voltage to always maintain the same dBm threshold, but this is unnecessary if we consider that the RSSI measurement has an accuracy of ±6dB. There is no point in wasting processing power to be precise, as the precision gain would be offset by the scarce accuracy of our hardware.

Another possibility for opportunistic transmissions is changing the modulation scheme according to the channel conditions. When the channel is very good, the use of higher-order modulation could trade power-efficiency for bandwidth-efficiency. This is of course not possible with current sensor network platforms (such as Mica2), where we do not have any flexibility at the physical layer (we are obviously constrained to one radio).

The mobile sink transmits at a fixed power level (0dBm), and the sensing nodes use this information along with the RSSI to adapt their transmission power. A simple look-up table maps ranges of RSSI to transmission power settings; a conservative approach has been adopted in the calibration of the look-up table to account for asymmetric links.

Figure 3.9 shows what happens in the sensing nodes each time they respond: they measure an RSSI value above the running average of all the RSSI values associated with all the packets they received. Node 1 is much closer to the sink than the others, so it can respond much more often than the others. The scales of the horizontal axes are different because not all nodes respond the same number of times. In this simple experiment, no packet loss was experienced. The transmission power control

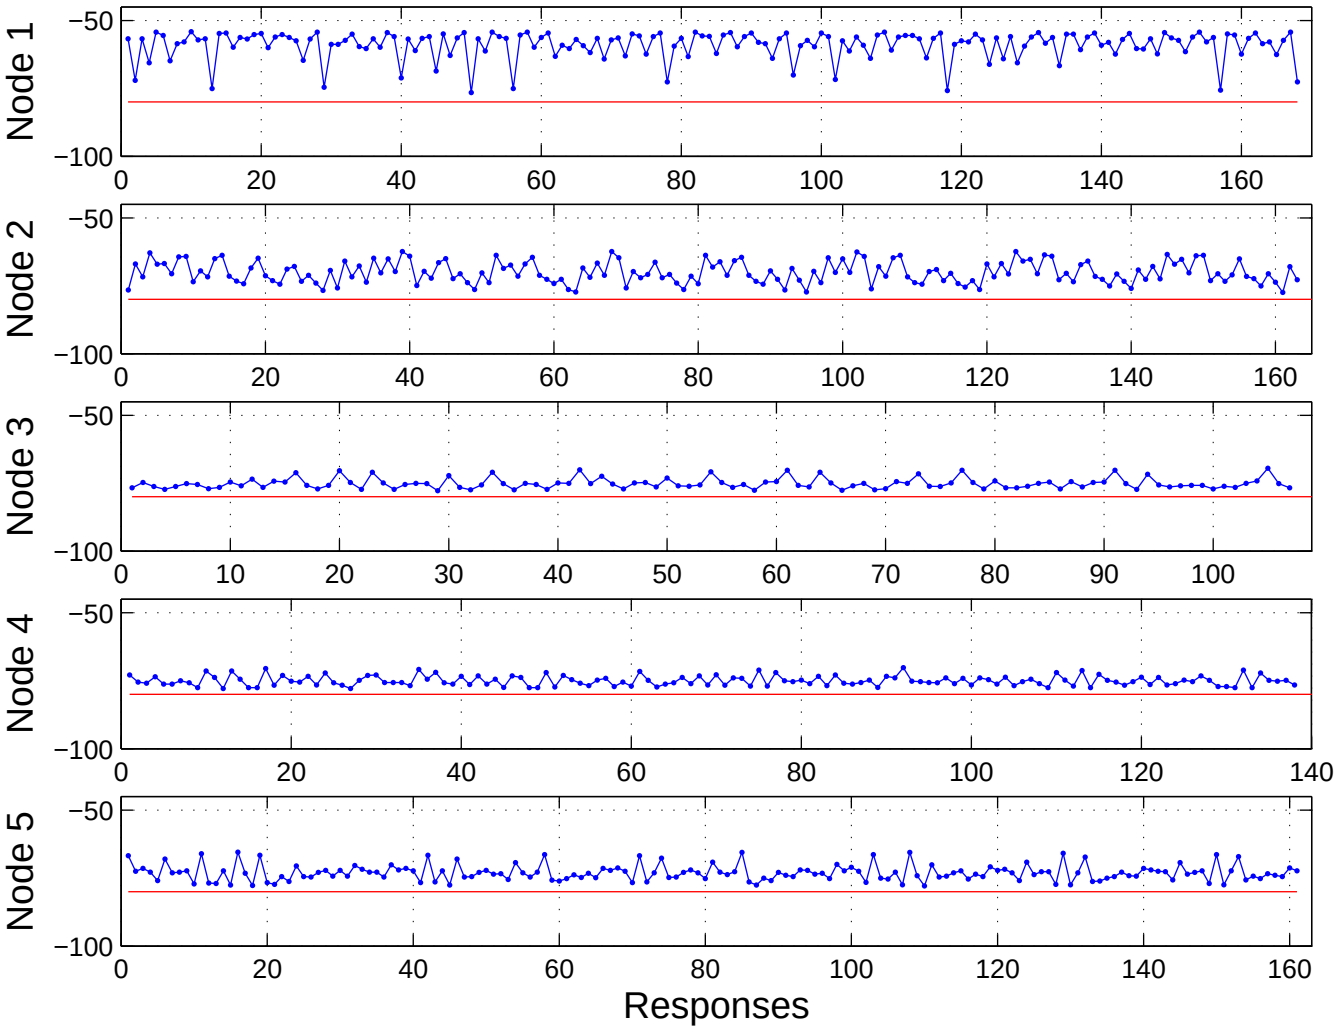

Figure 3.9. The sensing nodes only transmit when the RSSI (here shown in dBm on the vertical axis) exceeds a fixed threshold of -80dBm. For each response of each node, we show the RSSI level seen by the node and the running average measured by the node, used to regulate the transmission power.

TABLE 3.1

EFFICIENCY AND POWER GAIN IN MOBIHOP

Node	Mean RSSI [dBm]	Mean RSSI when responding [dBm]	Empirical Efficiency	Analytical Efficiency	Max. Power Gain [dB]
1	-60.4	-59	96%	99%	21
2	-73	-70.1	93%	94%	9.9
3	-80	-75.1	61%	73%	4.9
4	-78	-75	79%	82%	5
5	-75.5	-73	92%	89%	7

feature was seldom exploited due to the small variance around the mean received signal (the motion of the mobile sink is limited in scope). Only node 1, located in the proximity of the turntable, was able to choose the lowest power setting.

The extent to which transmission power control is used and the consequent power gain depends on the calibrations, but it is interesting to look at the highest possible power gain that can be achieved by MobiHop. Nodes respond to a request from the i-th block thereof if the measured RSSI r_i is above a threshold θ; if we assume the links to be perfectly symmetric, then in principle an upper bound for the power gain G_i defined as the difference between the fixed transmission power P_s of the mobile sink and the adaptable transmission power of the node P_n is given by

$$G_i \triangleq P_s - P_n = r_i - \theta \tag{3.6}$$

The quantity $\frac{G_i}{N}$, where N indicates the number of responses from a particular node, can be used as a figure of merit quantifying an upper bound to the power gain that can be achieved by the node.

The efficiency measured experimentally is remarkably close to the one predicted using results for Rayleigh fading channels. Efficiency is defined as the ratio between responses and requests: ideally, nodes should be able to respond to each request, but

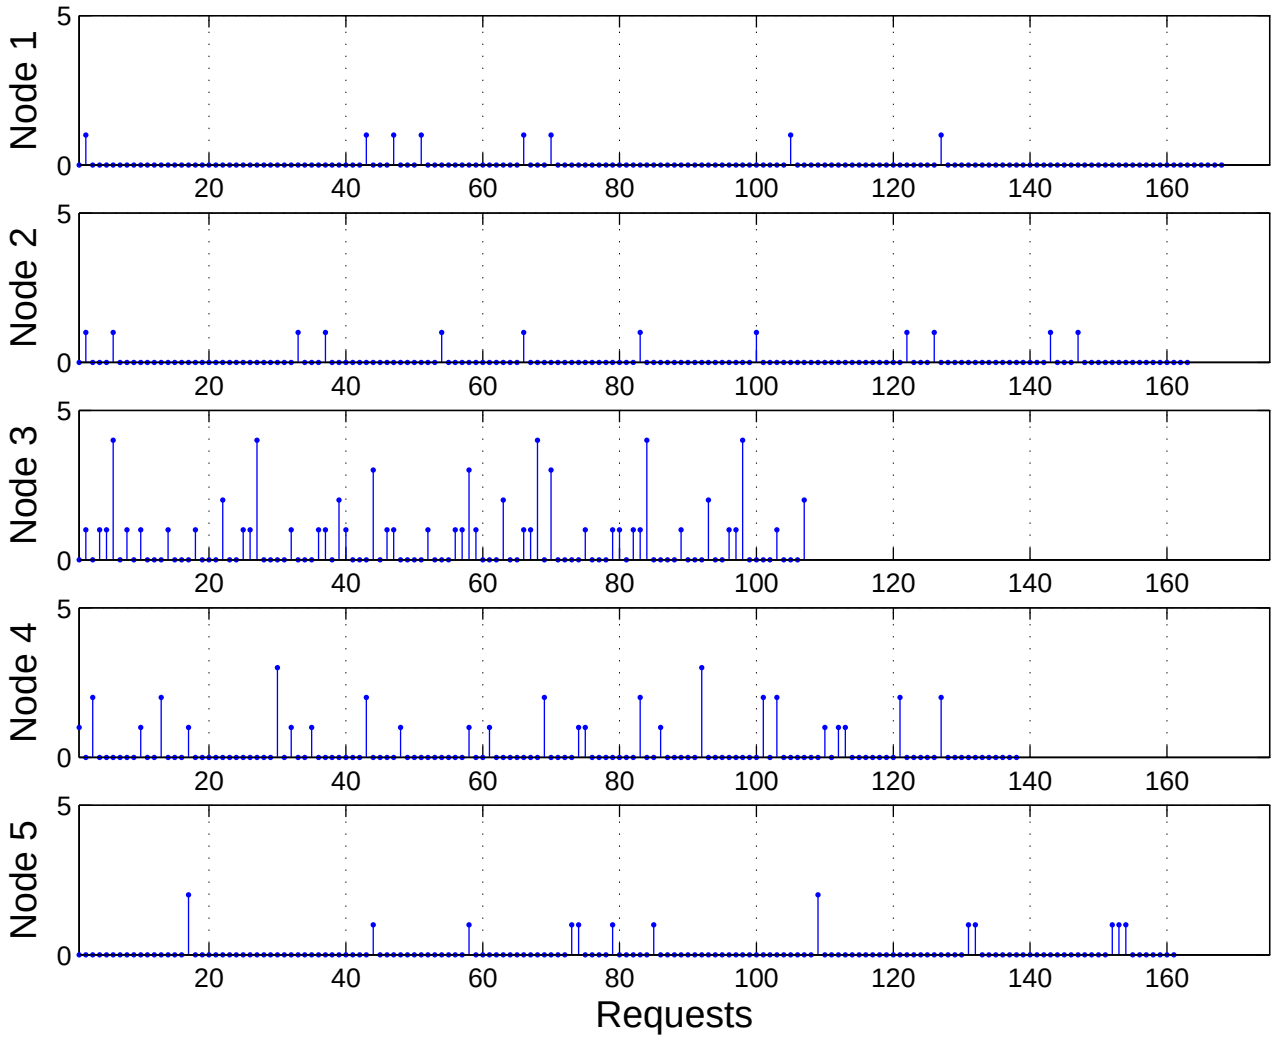

Figure 3.10. Owing to the opportunistic scheme applied by the sensing nodes, requests from the mobile sink are subjected to a variable delay. This Figure shows the delay for each request to each node. Delay is measured as the number of unanswered requests placed by the mobile sink.

in our scheme they may decide not to respond due to bad channel conditions. Our scheme provides a tradeoff between efficiency and energy retention (no wasted energy for useless transmissions resulting into packet loss). In Table 3.1, the column **Max. Power Gain** indicates the power gain upper bound for each node. The average power gain (over all the 5 nodes) is 9.6dB.

Figure 3.10 shows the delay suffered by each request. The delay is defined as the number of redundant requests needed for a node to respond to the mobile sink (number of unanswered requests). Large delays are rare due to the small number of nodes.

In practice, arbitrary motion is appealing: it can be achieved in a number of different ways and does not require any control or planning. Data mules are indeed being used in many existing sensor network deployments [11, 50]. In addition, arbitrary motion fits well into the MobiHop scheme, as it decorrelates the channels seen by a particular node at subsequent round robins.

With the deployment shown in Figure 3.8, we have added arbitrary motion in the form of a person walking around holding the mobile sink. Packet loss was not experienced, and the delay results are very good (the efficiency is greater than 95% for all nodes). This is because of the topology of the sensor network: the nodes are placed along the walls, and the arbitrary motion pattern always occurs within the limits defined by the nodes. In this example, the average upper bound for power gain is 14.1dB (16.9dB for node 1, 12.7dB for node 2, 13.5dB for node 3, 14.9dB for node 4, and 12.4dB for node 5).

We now look at another example of the behavior of MobiHop in a different indoor setting. We deployed the five sensing nodes in a hallway in the form of a line network, with nodes 10m apart from one another. We placed the mobile sink on a radio-controlled mini-vehicle and drove it back and forth alongside the line

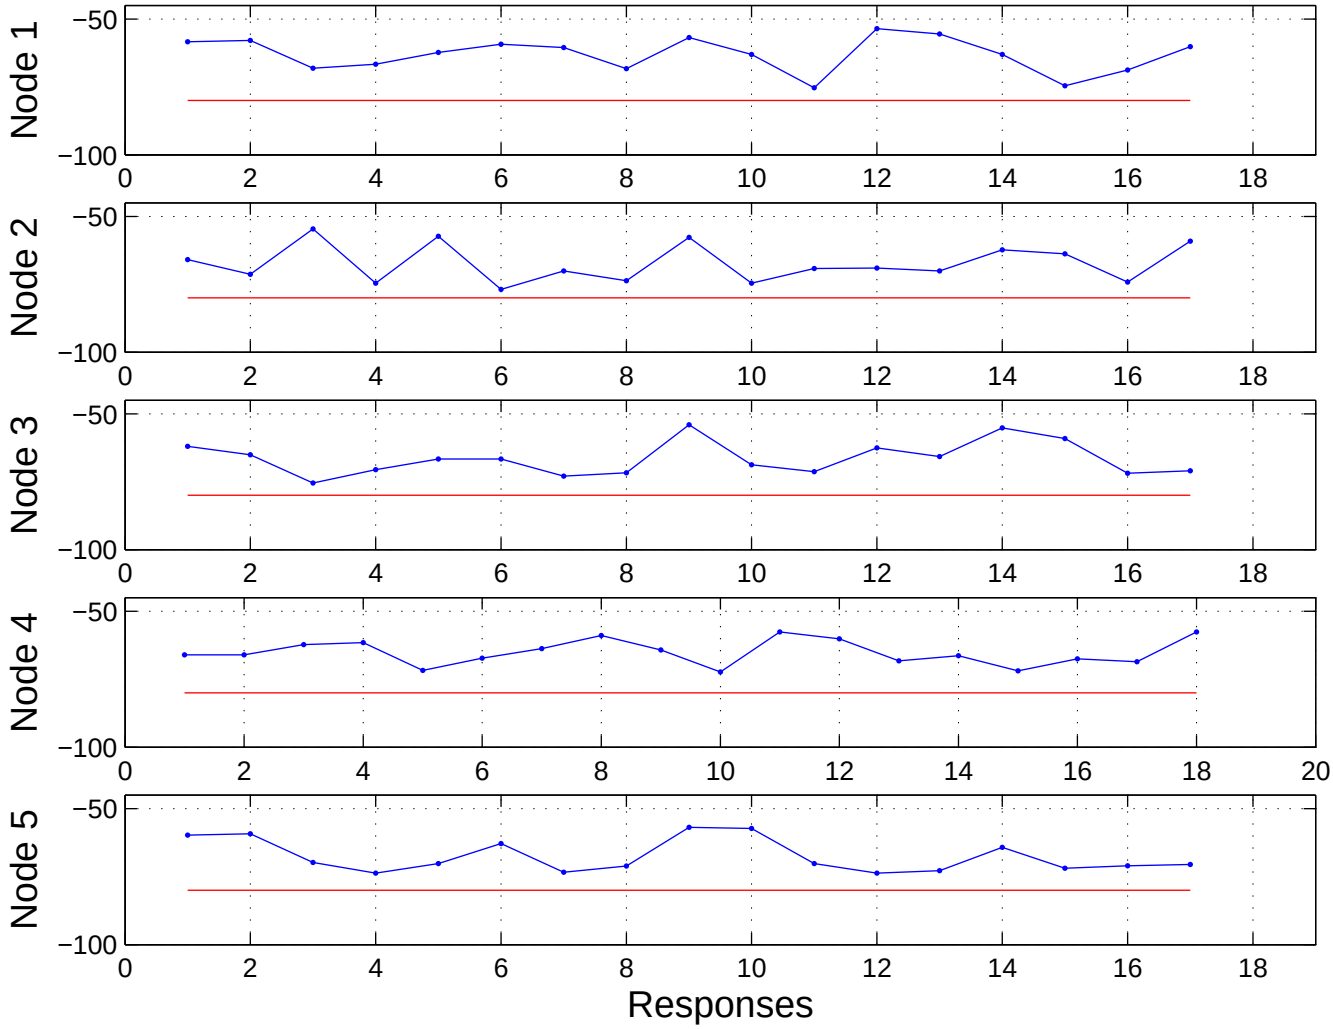

Figure 3.11. Arbitrary motion introduces a large variance in the RSSI values seen by the sensing nodes.

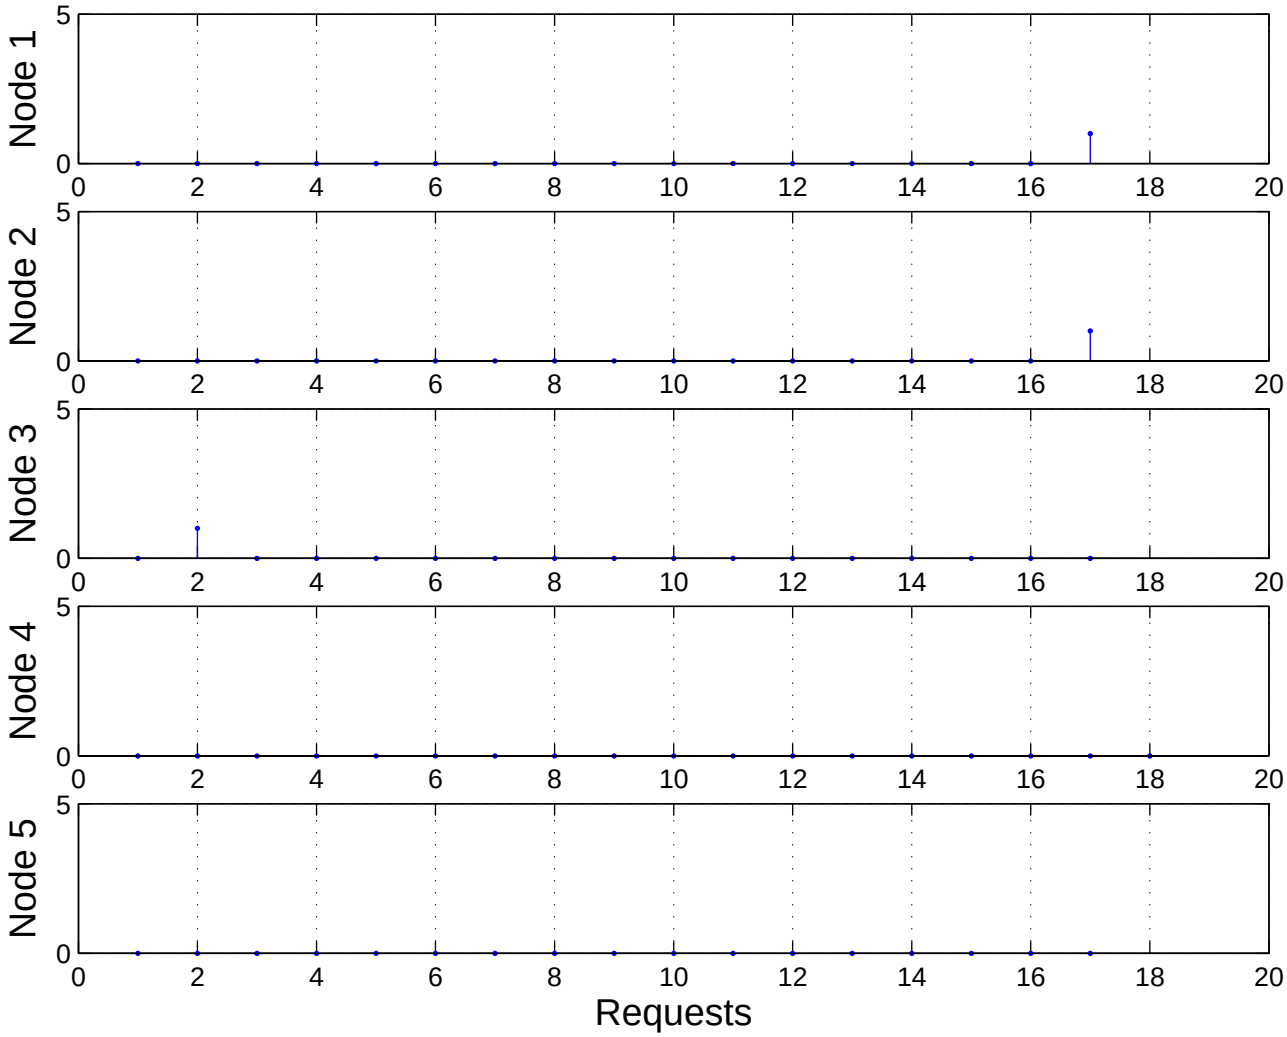

Figure 3.12. Arbitrary motion offers a higher degree of spatial diversity, which results in a lower delay.

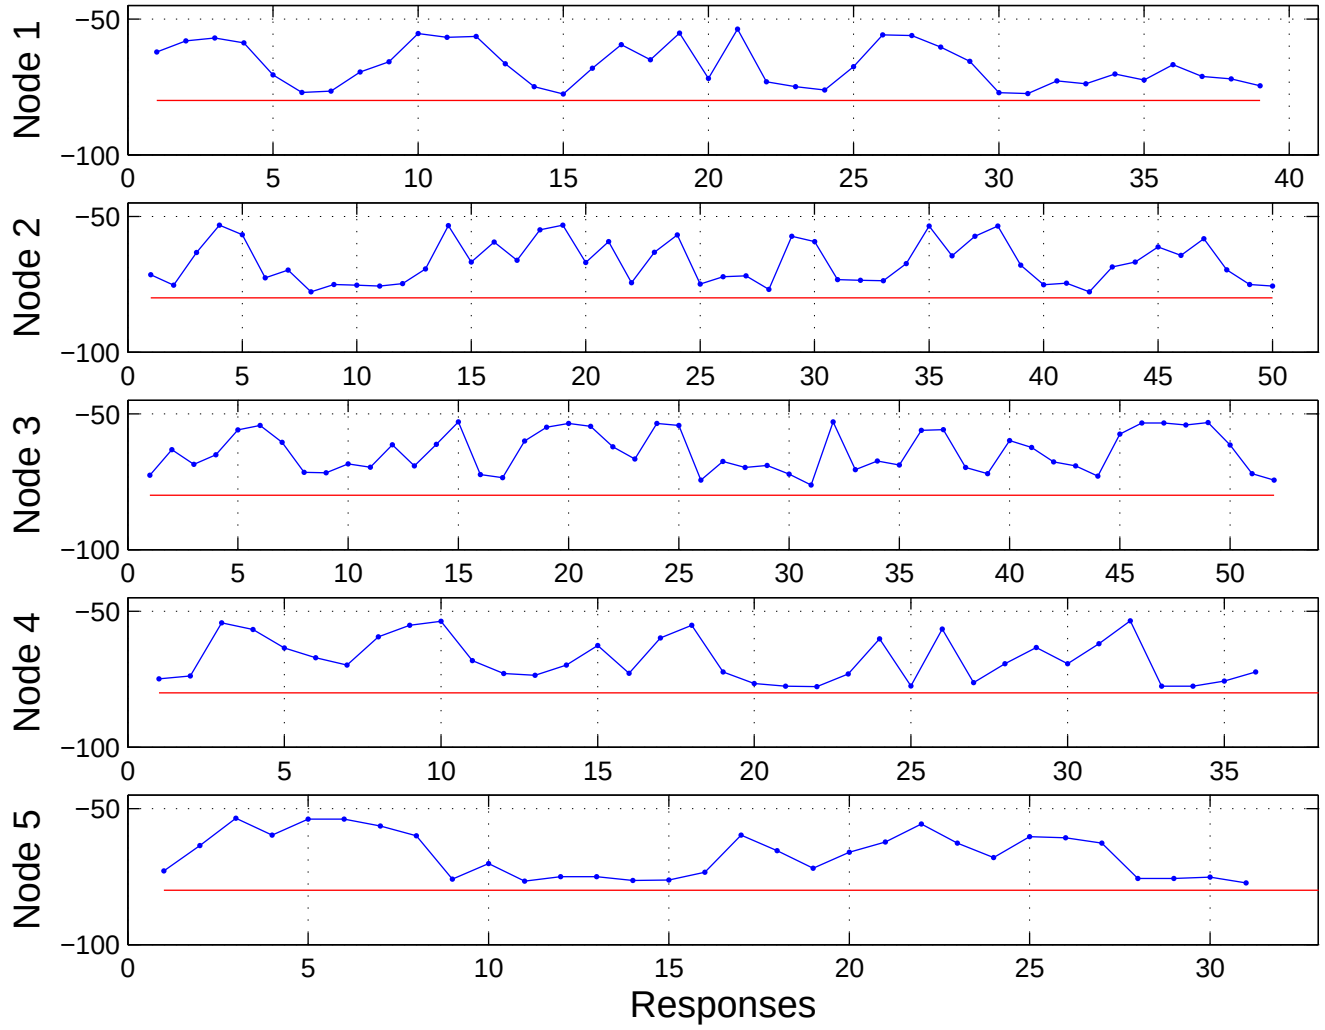

Figure 3.13. Arbitrary motion in a hallway with a line network topology. The sensing nodes reply only if the channel is better than a fixed threshold.

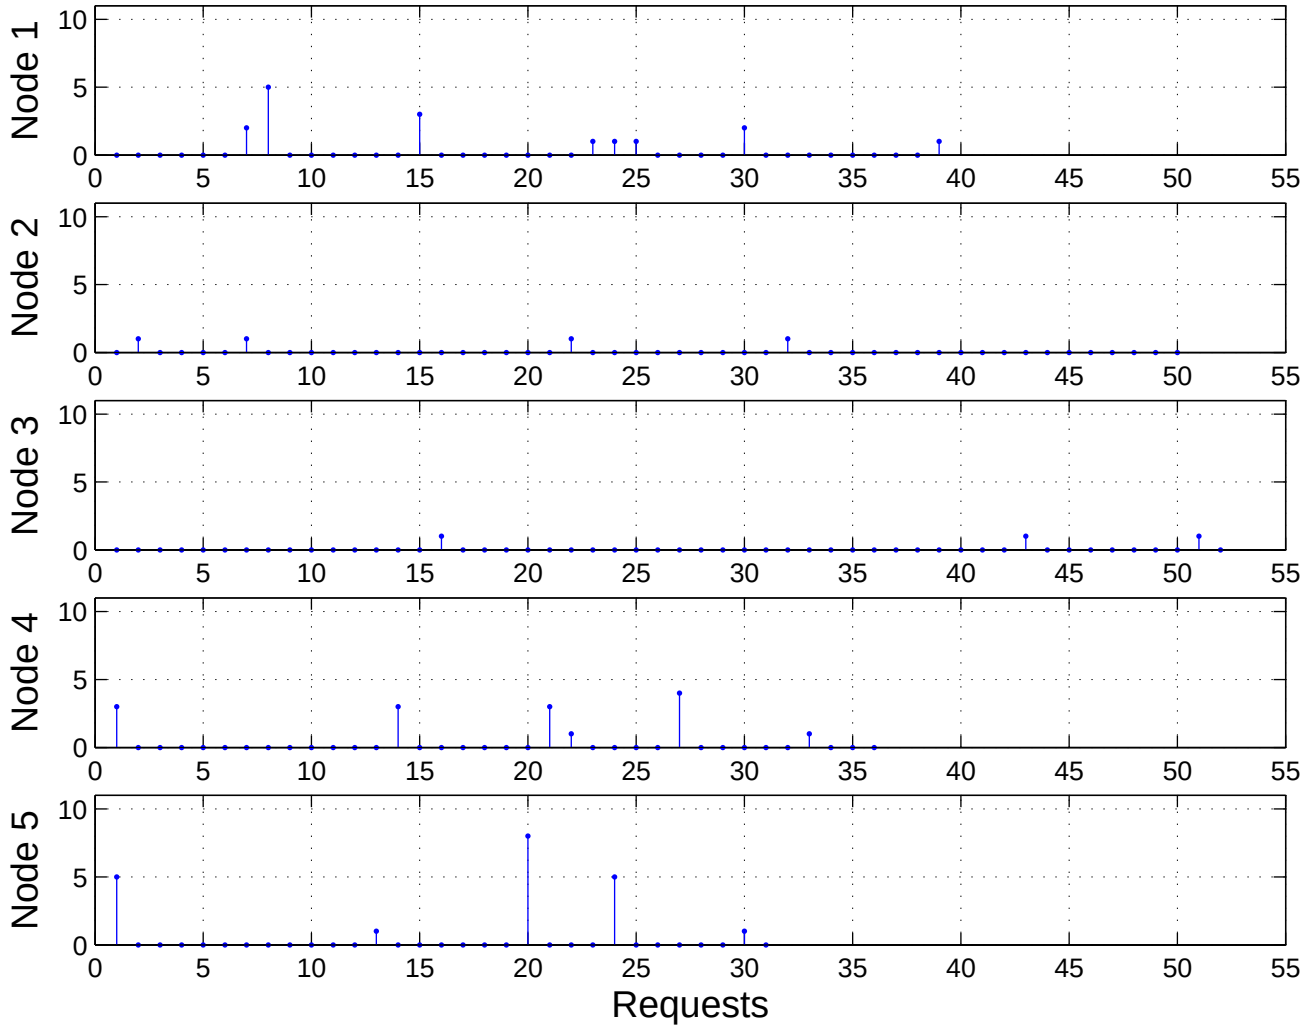

Figure 3.14. Arbitrary motion in a hallway with a line network topology.

of motes with an arbitrary motion pattern. The delay results (Figure 3.14) are very good if we consider the network topology and the absolutely arbitrary motion of the mobile sink. Arbitrary motion is indeed a worst-case scenario, and even so MobiHop exhibits no packet loss as well as reasonable delay results. In this example, the average upper bound for power gain is 13.5dB (13dB for node 1, 13dB for node 2, 12.4dB for node 3, 13.1dB for node 4, and 16dB for node 5). More details on the implementation can be found in the appendix.

3.4 Conclusions

We have designed and implemented MobiHop, a simple yet effective protocol for the energy-aware retrieval of data from a sensor network. MobiHop revolves around the concept of a mobile sink, a lightly constrained node bridging the base station and the sensing nodes. By moving around, the mobile sink makes the most out of spatial diversity and takes full advantage of the properties of the fading function. We have described various implementation possibilities along with the rationale behind each choice. We have shown how the Rayleigh fading model can be used to analyze the performance of MobiHop, and we have explored the behavior of our strategy with arbitrary motion. We believe that many sensor network application could use MobiHop to extend the lifetime of sensing nodes. The sensor network community is oriented toward lightweight multihop routing strategies, which have several drawbacks which are often overlooked. With MobiHop, we show how having a node move around even randomly can make single-hop data transfers possible minimizing the power consumption at the sensing nodes and preventing them from transmitting in the presence of bad channels. MobiHop is based on the use of RSSI as a link quality metric. The reliability of RSSI is maximized by selective snooping (which can be performed using low-power listening modi). Resilience to interference

can be implemented using our interesting novel idea for a protocol to distinguish noise from interference, which we have briefly mentioned here and whose thorough description will be the subject of future work.

CHAPTER 4

CONSTRUCTIVE EXPLOITATION OF STATIC MULTIPATH FADING: INFORMATION HARVESTING FROM SIGNAL STRENGTH MEASUREMENTS

The novel idea presented in this chapter demonstrates how any wireless network can be used as a wireless sensor network for motion detection without the addition of sensing hardware. The transceivers on wireless terminals are generally able to measure the strength of the received signal. Motion of individuals or objects in the vicinities of the terminals produces shadowing and multipath fading effects which alter the received signal strength and may therefore be detected. In this chapter we show how motion can be detected by analyzing variations in the signal strength indication, we present our experimental results, and we explore the potential impact of this powerful idea.

4.1 Motion Detection

Figures 4.1 and 4.2 show the outcome of simple point-to-point experiments using Mica2 motes. Figure 4.1 shows a variation in the signal strength measured by one node as the other is moved to a different position in the same room. It should be observed that a different signal strength is measured as the motion ends due to the new spatial coordinates of the pair. The same fading function associates a different channel gain to the new coordinates: a different set of coordinates corresponds to

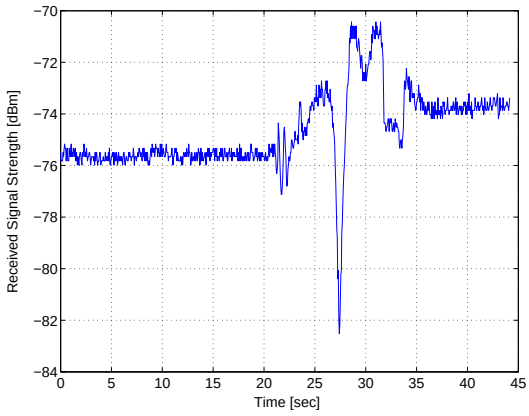

Figure 4.1. Effects of node motion

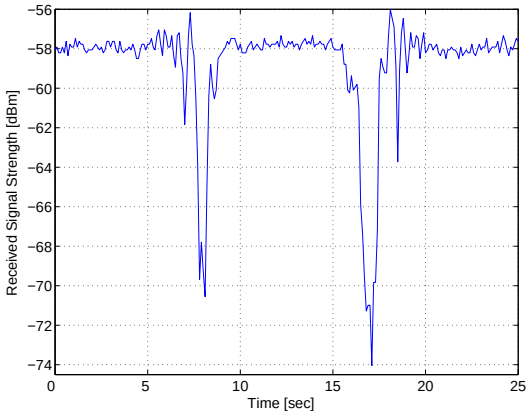

Figure 4.2. Effects of external motion

a different level of multipath fading. Figure 4.2 illustrates rapid changes in the measured signal strength produced by the passage of two individuals between the pair. The bodies of the individuals obstruct the line-of-sight path and cause a shadowing effect which is interpreted by the nodes as a topology change. The nodes observe different fading functions as motion is in progress and the individuals keep changing their position within the environment, until motion ends and the original fading function is restored. Motion of the terminals and motion of bodies other than the terminals (external motion) causes changes in the RSSI measured by the receiver mainly due to multipath fading. Shadowing is the dominant phenomenon when external motion occurs between the terminals. These observations form the basis for our implementation of a sensorless intrusion detector with a minimal network of three motes. A pair of nodes (sensing pair) is placed by the sides of the entrance to a room. A wireless link is established between them in the form of unidirectional packet flow (one acts as a transmitter, the other as a receiver). The receiver measures the strength of the signal from the transmitter and relays an array of readings to a base station connected to the external infrastructure via a gateway device. Application-level post-processing interprets the data from the nodes and notifies the user of intrusions. Various filtering strategies may be used; one example is the computation of the maximum difference between signal strength samples within a sliding window of M samples, as given by the equation

$$r_k = \max_{k-M<i\leq k} s_i - \min_{k-M<i\leq k} s_i, \tag{4.1}$$

where s is the array of M signal strength readings collected by the receiving node and r_k is the motion indicator computed at time k. The choice of M depends on the desired time window which can be configured by the receiving node on the basis of the transmission rate. An application running on the gateway or the external in-

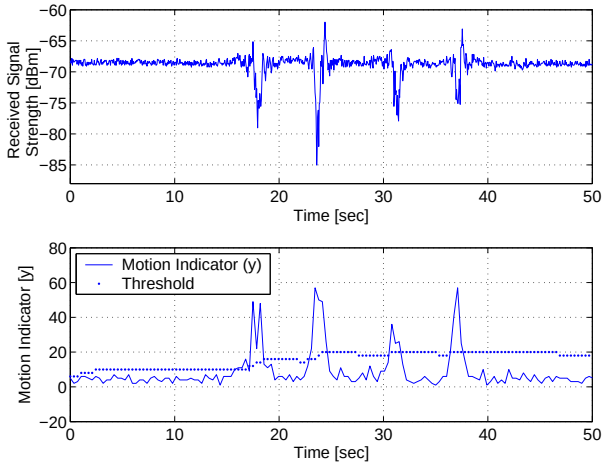

Figure 4.3. Sensorless intrusion detection

frastructure performs intrusion detection by comparing r_k to an adaptive threshold corresponding to a scaled version of the time average of r_k. This threshold is initialized to a value obtained with a learning stage during which the network acquires channel information: the application learns the channel gain to which the spatial coordinates of the sensing pair are mapped by the fading function. The user is to be notified of an intrusion whenever r_k exceeds the threshold. We have successfully implemented this idea with three Mica2 motes, one gateway device (MIB510 programming board), and a PC for post-processing. As shown in Figure 4.3, four individuals entering a room through a door monitored by the sensing pair cause noticeable variations in the measured signal strength. The base station computes r_k on the basis of the RSSI readings and compares it to an adaptive threshold. This concept can be extended to a larger network comprised of various pairs of opportunely placed nodes, with the caveat that interference between pairs should be avoided at the MAC layer. Nodes would need to keep track of the condition of the various links in which they participate.

4.2 Potential Impact

Our idea is particularly interesting as a tool to be applied in existing wireless networks, which are turned into sensor networks for motion detection at no additional cost. For instance, any 802.11-based wireless network can be used to log the signal strength in a particular environment. Fluctuations in the signal strength indicate motion in the surroundings. Users could simply leave their wireless network on and later refer to the signal strength log to extract motion-related information betraying a particular kind of activity. Parents coming back home after an evening out could see until what time their children took advantage of their absence, and burglary victims could determine when the intrusion took place. We have shown how nodes can be programmed to detect motion by means of a simple algorithm, but the variations in the received signal strength in the presence of intrusions shown in Figure 4.3 are so evident that they can be easily recognized by any wireless network user with no background in wireless communications. Other forms of information harvesting from the RSSI are also possible. One example is obtaining speed information from the observation of the distance between local extrema in the signal strength, which is related to the coherence time of the channel as observed by the receiver: a shorter coherence time corresponds to a higher speed.

We have introduced the exploitation of signal strength variations for purposes of sensorless motion detection in wireless networks. Our results, obtained without loss of generality in point-to-point settings, underline that multipath fading is a spatial phenomenon. The fading function captures a channel gain for a link between two nodes on the sole basis of the position of the nodes with respect to the topology of the environment. With a simple experimental setup, we decode the information contained in the fluctuations of the RF signal strength over a wireless link in order to detect motion in the vicinities of the nodes forming the link. This form of motion

detection is particularly appealing as an added feature to existing networks. It provides a minimal overhead, can be performed by existing sensor networks, and has a large number of interesting applications.

APPENDIX A

METHODOLOGY

Since the hardware of choice for our implementation of MobiRoute is Mica2, we used TinyOS components to implement our protocol. At a basic level, the interaction of TinyOS with the hardware may be described as follows: low-level components interface with the hardware, recognize interrupts and signal them to the higher-level components in the form of **events**. The higher-level components are thus able to recognize hardware interrupts in the form of events and can post computational **tasks** or call **commands** which are sent to the lower-level components.

In this section we intend to give some details on how MobiRoute was implemented in TinyOS. The mobile sink is a Mica2 mote with the software component MS installed (Figure A.1). After the MCU boots, writing to Flash is enabled and a timer is initialized to fire once a second. The timer controls the round robin: each time the timer fires, an **event** is signaled and a task is posted to send a request to the sensing node indicated by the polling counter. A separate counter is maintained for each node. The request is inserted in a packet with a broadcast address, so that nodes will not filter it out on the basis of the address. A port number (AM type) corresponding to the address of the intended destination is placed on the packet. If a packet is received from the radio, an **event** is signaled and another task is posted, where information is extracted from the received packet and the polling counter is incremented. In our implementation, we also write the extracted information to

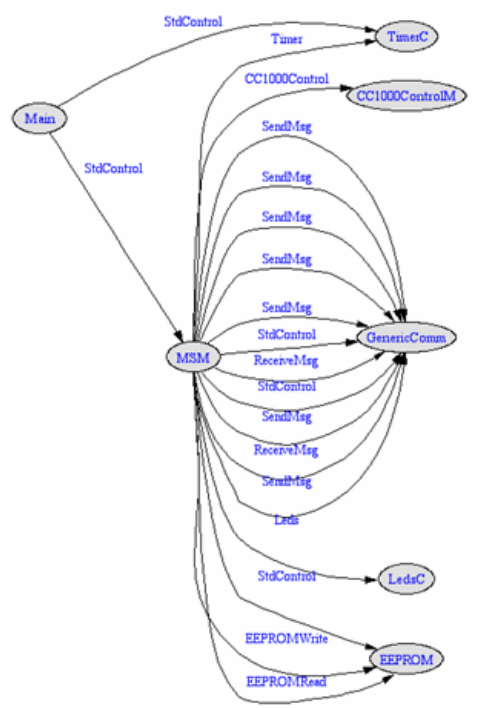

Figure A.1. The structure of the Mobile Sink (MS) TinyOS component

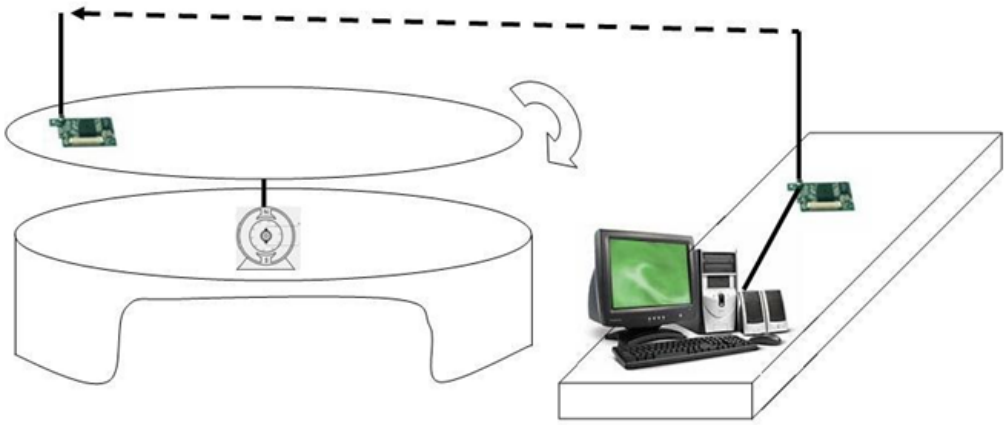

Figure A.2. The structure of the TinyOS sensing node component.

Flash and do not explicitly relay the packet to the base station. We gathered our results snooping packets with an external base station and used the Flash log to fill in the gaps to lost packets (packets received by the mobile sink which the base station could not overhear). Flash log retrieval is simply implemented with a UART interface which allows the mote to understand commands from the serial port. As a packet is received from the UART, the timer is masked to stop the polling, Flash writing is suspended, and a reading cycle begins. An LED warns the user of the ongoing log retrieval; as the LED is turned off, polling starts over.

In the sensing nodes, **events** are triggered as requests from the mobile sink are received. A task is posted for the computation of two metrics: a global average over

all RSSI samples measured since the MCU booted and an average over the latest 5 samples, which are stored in a FIFO buffer (buffer average). In addition, if the requests appear on the port whose number corresponds to the address of the node, a task for opportunistic transmissions is posted. If the signal level associated with the request is better than -80dB, the sensing node responds to the request from the sensing node. In our implementation, no actual sensor data are sent: the packet from the sensing nodes contain the RSSI and counter value associated to the request, and the value of the moving average used to choose the transmission power.

The moving average of the channel conditions is implemented as a linear combination of the buffer average and the global average:

$$\texttt{MovingAverage} = a\ \texttt{BufferAverage} + (1 - a)\ \texttt{GlobalAverage} \tag{A.1}$$

where a is a weighting factor that depends on the transmission rate of the mobile sink and its speed, and was set to 0.1 in our case.

BIBLIOGRAPHY

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

[2] Shooter Localization. `http://www.isis.vanderbilt.edu/projects/nest/applications.html`.

[3] The ZebraNet Wildlife Tracker. `http://www.princeton.edu/~mrm/zebranet.html`.

[4] Wireless Measurement and Control of the Indoor Environment in Buildings. `http://www.cbe.berkeley.edu/research/briefs/Wireless.htm`.

[5] I. F. Akyildiz, W. Su, Y. Sankarasubramaniam, and E. Cayirci. A Survey on Sensor Networks. In *IEEE Communications Magazine*, pages 102–114, August 2002.

[6] M. Beigl, A. Krohn, T. Zimmer, C. Decker, and Philip Robinson. AwareCon: Situation Aware Context Communication. In *The Fifth International Conference on Ubiquitous Computing (Ubicomp'03)*, Seattle, WA, October 2003.

[7] S. Bergbreiter and K. S. J. Pister. CotsBots: An Off-the-Shelf Platform for Distributed Robotics. In *Proceedings of the 2003 IEEE International Conference on Intelligent Robots and Systems (ICRA'03)*, Las Vegas, NV, October 2003.

[8] J. Beutel, O. Kasten, F. Mattern, K. Roemer, F. Siegemund, and L. Thiele. Prototyping Sensor Network Applications with BTnodes. In *IEEE European Workshop on Wireless Sensor Networks (EWSN'04)*, Berlin, Germany, January 2004.

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

[10] E. Biagioni and K. Bridges. The Application of Remote Sensor Technology to Assist the Recovery of Rare and Endangered Species. *International Journal of High Performance Computing Applications*, 16(3):315–324, August 2002.

[11] J. Burrell, T. Brooke, and R. Beckwith. Vineyard Computing: Sensor Networks in Agricultural Production. *IEEE Pervasive Computing*, 3(1):38–45, 2004.

[12] A. Cerpa, J. Elson, D. Estrin, L. Girod, M. Hamilton, and J. Zhao. Habitat Monitoring: Application Driver for Wireless Communications Technology. In *ACM SIGCOMM Workshop on Data Communications in Latin America and the Caribbean*, San Jos, Costa Rica, April 2001.

[13] A. Chakrabarti, A. Sabharwal, and B. Aazhang. Using Predictable Observer Mobility for Power Efficient Design of Sensor Networks. In *Information Processing in Sensor Networks (IPSN'03)*, Palo Alto, CA, April 2003.

[14] A. P. Chandrakasan, R. Min, M. Bhardwaj, S. Cho, and A. Wang. Power Aware Wireless Microsensor Systems. In *28th European Solid-State Circuits Conference (ESSCIRC'02)*, Florence, Italy, 2002.

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

[16] T. Fulford-Jones, D. Malan, M. Welsh, and S. Moulton. CodeBlue: An Ad Hoc Sensor Network Infrastructure for Emergency Medical Care. In *International Workshop on Wearable and Implantable Body Sensor Networks*, London, UK, 2004.

[17] D. Ganesan, R. Govindan, S. Shenker, and D. Estrin. Highly Resilient, Energy Efficient Multipath Routing in Wireless Sensor Networks. In *Proceedings of the 2nd ACM International Symposium on Mobile Ad Hoc Networking and Computing (MobiHoc'01)*, pages 251–254, Long Beach, CA, 2001.

[18] P.B. Gibbons, B. Karp, Y. Ke, S. Nath, and S. Seshan. IrisNet: An Architecture for a Worldwide Sensor Web. *IEEE Pervasive Computing*, 2(4):22–33, 2003.

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

[20] A. Goldsmith and S. Wicker. Design challenges for energy-constrained ad hoc wireless networks . *IEEE Wireless Communications Magazine*, 9(4):8–27, August 2002.

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

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

[23] M. Haenggi. On Routing in Random Rayleigh Fading Networks. *IEEE Transactions on Wireless Communications*, 2004. Accepted for publication. Available at `http://www.nd.edu/~mhaenggi/pubs/routing.pdf`.

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

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

[26] W. C. Jakes. *Microwave Mobile Communications*. IEEE Press, New York, NY, 1974.

[27] A. Kansal, A. A. Somasundara, D. D. Jea, M. B. Srivastava, and D. Estrin. Intelligent fluid infrastructure for embedded networks. In *Proceedings of the 2nd International Conference On Mobile Systems, Applications And Services*, pages 111–124, Boston, MA, 2004.

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

[29] K. Van Laerhoven, N. Villar, and H.-W. Gellersen. Pin&Mix: When Pins Become Interaction Components. In *Physical Interaction (PI03) Workshop on Real World User Interfaces" - Mobile HCI Conference*, Udine, Italy, September 2003.

[30] D. Lal, A. Manjeshwar, F. Herrmann, E. Uysal-Biyikoglu, and A. Keshavarzian. Measurement and Characterization of Link Quality Metrics in Energy Constrained Wireless Sensor Networks. In *Proceedings of the IEEE 2003 Global Communications Conference (GLOBECOM'03)*, pages 446–452, San Francisco, CA, December 2003.

[31] J. Lifton, D. Seetharam, M. Broxton, and J. Paradiso. Pushpin Computing System Overview: A Platform for Distributed, Embedded, Ubiquitous Sensor Networks. In *Proceedings of the Pervasive Computing Conference*, Zurich, Switzerland, August 2002.

[32] A. Mainwaring, J. Polastre, R. Szewczyk, D. Culler, and J. Anderson. Wireless Sensor Networks for Habitat Monitoring. In *First ACM Workshop on Wireless Sensor Networks and Applications*, Atlanta, GA, September 2002.

[33] T. Martin, M. Hsiao, D. Ha, and J. Krishnaswami. Denial-of-Service Attacks on Battery-powered Mobile Computers. In *Proceedings of the 2nd IEEE Pervasive Computing Conference*, pages 309–318, Orlando, FL, March 2004.

[34] M. Mauve, H. Hartenstein, H. Fuessler, J. Widmer, and Wolfgang Effelsberg. Positionsbasiertes Routing fuer die Kommunikation zwischen Fahrzeugen. *it Information Technology (formerly it+ti) - Methoden und innovative Anwendungen der Informatik und Informationstechnik*, 44:278–286, October 2002.

[35] M. B. McMickell, B. Goodwine, and L. A. Montestruque. MICAbot: a robotic platform for large-scale distributed robotics. In *Proceedings of International Conference on Intelligent Robots and Systems (ICRA'03)*, volume 2, pages 1600–1605, Taipei, Taiwan, 2003.

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

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

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

[39] J. A. Paradiso, J. Lifton, and M. Broxton. Sensate media multimodal electronic skins as dense sensor networks. *BT Technology Journal*, 22:32–44, October 2004.

[40] S. Park, I. Locher, and M. Srivastava. Design of a Wearable Sensor Badge for Smart Kindergarten. In *6th International Symposium on Wearable Computers (ISWC2002)*, pages 13.1–13.22, Seattle, WA, October 2002.

[41] J. Polastre, R. Szewczyk, and D. Culler. Telos: Enabling Ultra-Low Power Wireless Research. In *The Fourth International Conference on Information Processing in Sensor Networks: Special track on Platform Tools and Design Methods for Network Embedded Sensors (IPSN/SPOTS)*, Los Angeles, CA, April 2005.

[42] T. S. Rappaport. *Wireless Communications*. Prentice Hall, Upper Saddle River, NJ, 1996.

[43] N. Reijers and K. Loangendoen. Efficient Code Distribution in Wireless Sensor Networks. In *Second ACM International Workshop on Wireless Sensor Networks and Applications*, San Diego, CA, September 2003.

[44] S. Roundy, P.K. Wright, and J. Rabaey. Simulation of flat fading using MATLAB for classroom instruction. *IEEE Transactions on Education*, 45:19–25, February 2002.

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

[46] J. Sallai, G. Balogh, M. Marti, and A. Ldeczi. Acoustic Ranging in Resource Constrained Sensor Networks. Technical Report ISIS-04-504, Vanderbilt University, 2004.

[47] A. Savvides and M. B. Srivastava. A Distributed Computation Platform for Wireless Embedded Sensing. In *20th International Conference on Computer Design (ICCD'02)*, Freiburg, Germany, September 2002.

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

[49] C. Schurgers, O. Aberthorne, and M. Srivastava. Modulation scaling for Energy Aware Communication Systems. In *Proceedings of the 2001 International Symposium on Low Power Electronics and Design*, pages 96–99, Huntington Beach, CA, August 2001.

[50] R. C. Shah, S. Roy, S. Jain, and W. Brunette. Data MULEs: Modeling and analysis of a three-tier architecture for sparse sensor networks. In *Ad Hoc Networks Journal*, volume 1, pages 215–233. Elsevier, Sep 2003.

[51] E. Shih, P. Bahl, and M.J. Sinclair. Wake on wireless: an event driven energy saving strategy for battery operated devices. In *Proceedings of the 8th Annual International Conference on Mobile Computing and Networking (MobiCom'02)*, Atlanta, GA, September 2002.

[52] F. Sivrikaya and B. Yener. Time synchronization in sensor networks: a survey. *IEEE Network*, 18(4):45–50, July–August 2004.

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

[54] F. Stajano and R. Anderson. The Resurrecting Duckling: Security Issues for Ad-hoc Wireless Networks. In *7th International Workshop on Security Protocols*, Cambridge, UK, April 1999.

[55] M. Strohbach. The Smart-its Platform for Embedded Context-Aware Systems. In *Proceedings of the First International Workshop on Wearable and Implantable Body Sensor Networks*, London, UK, April 2004.

[56] L. F. W. van Hoesel, S. O. Dulman, P. J. M. Havinga, and H. J. Kip. Design of a low-power testbed for Wireless Sensor Networks and verification. Technical Report R-CTIT-03-45, University of Twente, September 2003.

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

[58] G. Werner-Allen, M. Welsh, J. Johnson, and M. Ruiz. Monitoring Volcanic Eruptions with a Wireless Sensor Network . Available at `http://www.eecs.harvard.edu/~werner/projects/volcano`.

[59] A. Woo, T. Tong, and D. Culler. Taming the Underlying Challenges of Reliable Multihop Routing in Sensor Networks. In *Proceedings of the 1st International Conference on Embedded Networked Sensor Systems (SenSys'03)*, Los Angeles, CA, November 2002.

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

[61] O. Younis and S. Fahmy. HEED: A Hybrid, Energy-Efficient, Distributed Clustering Approach for Ad-hoc Sensor Networks. In *IEEE Transactions on Mobile Computing*, volume 3, pages 366–379, 2004.

[62] L. Yuan and G. Qu. Energy-Efficient Design of Distributed Sensor Networks. In M. Ilyas and I. Mahgoub, editors, *Handbook of Sensor Networks: Compact Wireless and Wired Sensing Systems*, pages 38.1–38.19, Boca Raton, FL, 2004. CRC Press.

[63] M. Zuniga and B. Krishnamachari. Analyzing the Transitional Region in Low Power Wireless Links. In *Proceedings of the First IEEE International Conference on Sensor and Ad hoc Communications and Networks (SECON'04)*, Santa Clara, CA, October 2004.