Eidgenössische Technische Hochschule Zürich
Swiss Federal Institute of Technology Zurich

Ultra-low Power Wireless Hierarchical Sensing

Master Thesis

Felix Sutton
fsutton@student.ethz.ch

Computer Engineering and Networks Laboratory
Department of Information Technology and Electrical Engineering
ETH Zürich

Supervisors:
Dr. Jan Beutel
Bernhard Buchli
Prof. Dr. Lothar Thiele

November 12, 2012

Acknowledgements

I would like to thank Prof. Dr. Lothar Thiele for giving me the opportunity to write this Master thesis in his research group.

I am very grateful to my supervisors, Dr. Jan Beutel and Bernhard Buchli, for their support and guidance throughout this exciting and challenging Master thesis. They empowered me to drive the technical direction of the project and push the limits of my knowledge.

I would also like to extend my gratitude to Dr. Thomas Burger, Thomas Kleier, Philipp Schönle (Integrated Systems Laboratory, ETH Zürich), Mustafa Boyvat, Marco Zahner (Electromagnetic Fields and Microwave Electronics Laboratory, ETH Zürich), and Dr. Stephan Gruber (Department of Geography, University of Zürich) for their expert advice.

Furthermore, I would like to take this opportunity to thank my family and my girlfriend for their amazing support throughout my studies.

Abstract

Wireless sensor networks have the potential to provide wireless data retrieval of environmental parameters, that would otherwise be too costly or prohibitive to acquire through manual campaigns. These traditional mote-based networks utilise wired sensors, which suffer from restricted sensor placement and limited sensing reliability.

An extension to a traditional mote-based sensor network, termed wireless hierarchical sensing, is introduced in this thesis. Using a sensor abstraction coupled with an energy efficient wake-up radio architecture, the new wireless hierarchical sensor network "extends the sensor wires" to provide flexible sensor placement and increased sensing reliability.

A realistic use case of ground surface temperature monitoring is used throughout this thesis to demonstrate the application of wireless hierarchical sensing. A functional prototype of a wireless hierarchical sensor network is developed, consisting of an ultra-low power cluster node, and a stand-alone cluster-head. The functional requirements, detailed system design, prototype implementation and system evaluation are presented. The cluster node achieves an average power consumption of 78.5μW in a realistic deployment configuration, leading to an estimated operation lifetime of more than 3 years from a single coin cell battery.

Contents

Chapter 1

Introduction

1.1 Motivation

Understanding the factors impacting slope stability in alpine regions continues to be an important area of research in the Geoscientific community. The seasonal melting of snow is one such factor which has been shown to contribute to the thawing of permafrost, and thus adversely affects slope stability in high mountain regions [61].

A recent study [46] investigated the small and large-scale topographical effects of ground surface temperature in the Swiss Alps. The study monitored ground surface temperature using commercial miniature temperature data loggers, called iButtons [84]. The alpine field site was strategically segmented into footprints of 10 x 10m in size, spanning a wide-range of topographical features. Each footprint contained 10 randomly positioned iButtons buried at a depth of 50mm below the Earth's surface. Each iButton was configured to record a temperature measurement every 3 hours.

The low-cost iButton data loggers enabled flexible sensor placement with high spatial and temporal coverage. However, the manual campaign to retrieve the data loggers in order to download the recorded temperature measurements was costly and prohibitive. Retrieving the buried iButton's accumulated excessive costs (i.e. in labour, transportation and accommodation for several days), and was hampered by volatile weather conditions and the remote location of the field sites.

Recent advances in wireless sensor technology can help remedy the challenges imposed by manual data retrieval campaigns [68]. As has been exemplified in the X-Sense [72] project, wireless sensor networks have been successfully deployed and operated in the Swiss Alps for long-term environmental monitoring [32]. The X-Sense wireless sensor networks provide a robust wireless data retrieval and back-end processing platform, thus eliminating costly manual data retrieval campaigns.

However, traditional mote-based wireless sensor networks impose restrictions on sensor placement and adversely impact sensing reliability. For example, the X-Sense PermaDAQ sensor node [31] multiplexes several single-channel wired sensors, aggregates the measurements and transmits the data over a multi-hop wireless sensor network, as depicted in Figure 1.1(left). The placement of the sensors requires laying cables between each sensor and the wireless sensor node. The sensor cabling is not only cumbersome to install, but also prone to physical damage. Furthermore, traditional wired sensing suffers from a single point of failure with respect to sensing capability. Since all wired sensors are directly controlled by the wireless sensor node, if the node becomes unavailable (e.g. due to a depleted power source), all sensing capability assigned to that node is lost.

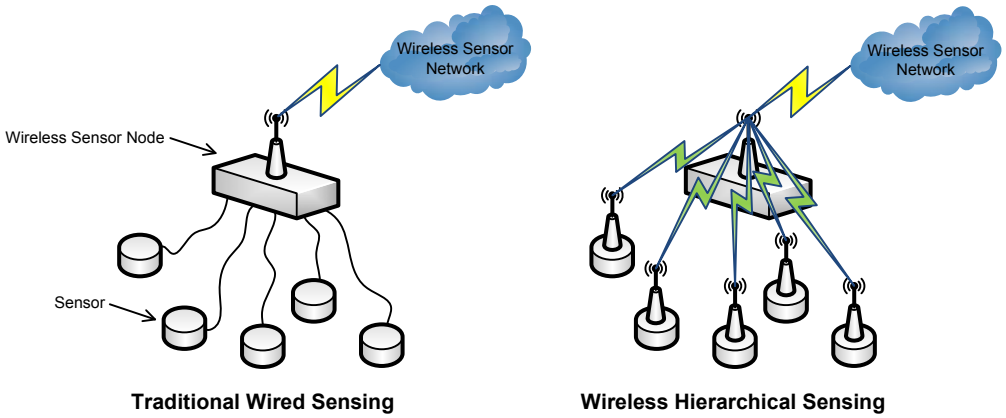

Figure 1.1: (left) Traditional wired sensing compared to (right) wireless hierarchical sensing.

An alternative to traditional wired sensing is to "extend the sensing wires" using a second-tier of wireless connectivity, thus forming a wireless hierarchical sensor network. Each sensor collects measurements autonomously, while also monitoring a point-to-point wireless link toward the wireless sensor node, as illustrated in Figure 1.1(right). The wireless sensor node requests measurement data from the surrounding sensors using an energy efficient wake-up radio architecture. The introduction of wireless hierarchical sensing decouples traditional mote-based sensing in both space and time. The wireless connectivity provides greater flexibility in the sensor placement (i.e. space dimension), while each sensor operates autonomously (i.e. time dimension) to improve sensing reliability.

This thesis introduces the new concept of wireless hierarchical sensing. A functional prototype of a wireless hierarchical sensor network is characterised and developed, suitable for later integration into a wireless sensor network. A real-world use case of ground surface temperature monitoring is used throughout this thesis to demonstrate the application of wireless hierarchical sensing.

1.2 Contributions

The contributions of this master thesis are summarised as follows:

- Requirement analysis of a wireless hierarchical sensor network, using the ground surface temperature sensing use case to demonstrate the network's functionality. *(Chapter 2)*
- Design and implementation of a prototype wireless hierarchical sensor network, consisting of an ultra-low power cluster node, and a stand-alone cluster-head node. *(Chapter 3)*
- Performance analysis of the functional prototype wireless hierarchical sensor network. *(Chapter 4)*

1.3 Background

1.3.1 Wireless Hierarchical Sensing

The aim of wireless hierarchical sensing is to provide flexible sensor placement and increased sensing reliability of traditional mote-based wireless sensor networks. This is achieved by extending the mote-based network with an additional level of wireless connectivity based on an energy efficient wake-up radio architecture. Together with an abstract representation of a sensor, one can effectively "extend the sensor wires" of a single sensor interface into a cluster of single-channel sensors. Each member in the cluster operates autonomously; collecting samples from a single-channel sensor and responding to data requests from the cluster-head over the wireless link. The wireless connectivity provides sensor placement flexibility, while the autonomous operation of each cluster member increases sensing reliability.

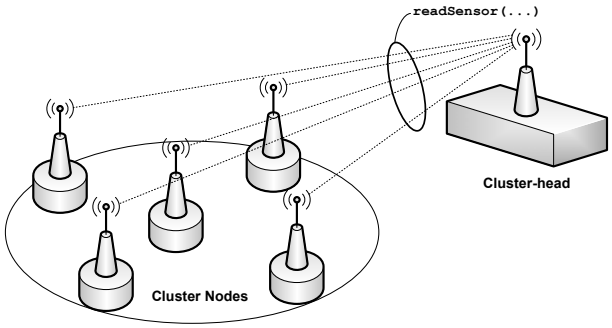

Figure 1.2: Architecture of a wireless hierarchical sensor network.

The architecture of a wireless hierarchical sensor network is illustrated in Figure 1.2. It consists of a cluster-head that requests data from a cluster of nodes through an abstract sensor function call (i.e. *readSensor(...)*). The cluster nodes utilise a wake-up radio (see Section 1.3.2) to handle the sensor data request in an energy efficient manner. Each cluster node responds to the cluster-head with the requested sensor data using a data radio coupled with a multi-user access policy. The generality of the sensor abstraction supports any sensor type and modality.

The cluster-head does not perform any form of topology control on the cluster (i.e. there is no clustering algorithm [28] executing on the cluster-head). Similarly, each cluster node has no knowledge of neighbouring devices within the cluster. It is assumed however, that each cluster node has a pre-assigned identity unique within its cluster, and that the cluster-head has sufficient power resources to request and receive measurement data from the cluster at a rate determined by the target application.

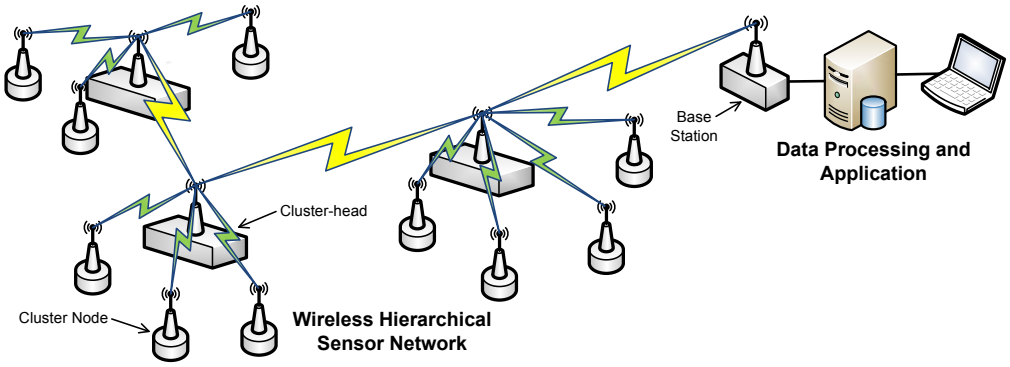

Figure 1.3: System concept of a wireless hierarchical sensor network.

Figure 1.3 illustrates the system concept of a wireless hierarchical sensor network for the purpose of ground surface temperature monitoring. A collection of ultra-low power cluster nodes is distributed in the vicinity of each cluster-head. The cluster nodes are configured to periodically sample a built-in temperature sensor and record the measurement and the time of acquisition. Each cluster-head is responsible to retrieve the measurement data from its nearby cluster. This is facilitated by broadcasting a wake-up sequence to all cluster nodes, which in turn triggers the nodes to awake from a low power mode and transmit the recorded temperature measurements using a Time Division Multiple Access (TDMA) policy.

The cluster nodes achieve end-to-end wireless data retrieval by using the network services available at the cluster-head. It is envisaged that each cluster-head executes an energy efficient network protocol stack. To that end, each cluster-

head aggregates the temperature sensor data from the cluster and forwards it to a base station over a multi-hop network. The base station is interfaced to a data processing infrastructure, where the received temperature data is stored and made available for the target application. For example, the collected temperature measurements can be used to investigate the dependence of ground surface temperature on topography [46] or be used to derive melting characteristics of seasonal snow cover [61].

1.3.2 Wake-up Radio Architecture

Radio communication is one of the most energy consuming tasks performed in wireless sensor networks. There exist several methods for reducing this heavy energy burden, with the most popular being the implementation of energy efficient MAC layer protocols (see [38] for a survey). The primary goal of such MAC layer protocols is to aggressively duty-cycle the radio, while still maintaining a minimum level of network connectivity.

An alternative method of reducing radio energy consumption is to approach the problem at the physical layer, whereby separate radio hardware is used for data transmission and channel monitoring functions [37]. The central idea is to idle listen on a wake-up channel with an ultra-low power radio, and only turn on the high power radio for data communication when indicated on the wake-up channel. Significant energy is saved as the high power data radio is only active during the necessary data communication.

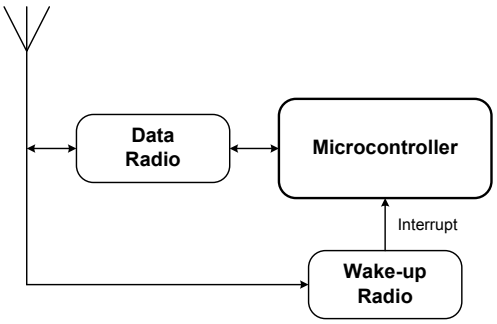

Figure 1.4: High-level block diagram of a wake-up radio architecture for wireless sensor nodes.

The high-level block diagram of architecture of a wireless sensor node supporting a wake-up radio is shown in Figure 1.4. The wireless sensor consists of a microcontroller for general processing, and two radio modules. The wake-up radio is an ultra-low power radio receiver which is continually searching for a specific wake-up sequence on the wake-up channel. If a wake-up sequence is suc-

cessfully detected, an interrupt is issued to the microcontroller, triggering it to turn on the data radio and commence communication.

A limitation of the wake-up radio architecture is its low-range and susceptibility to false wake-ups. The communication range of a wake-up radio architecture is determined by the minimum range of the wake-up and data radios. Since the wake-up radio is designed for ultra-low power dissipation, it typically suffers from a lower receiver sensitivity relative to the high power data radio. The asymmetric receiver sensitivity restricts the total communication range to that of the wake-up radio reception range. The low wake-up receiver sensitivity also makes it possible for noise or interference to erroneously trigger a wake-up.

There are benefits, challenges and trade-offs associated with wake-up radios, as summarised in [39]. However, there exist several application areas where the benefits of a wake-up radio outweigh all other pitfalls. One such example is its use in wireless hierarchical sensor networking introduced in this thesis.

1.3.3 Ultra-low Power System Design

There is no consensus in the literature on the exact definition of an ultra-low power system, as one can interpret the constraints implied on the system in a multitude of ways. However, there is one rather holistic approach which seeks to capture the fundamental design trade-offs imposed by such a system. As represented diagrammatically in Figure 1.5, one may view an ultra-low power system as a complex trade-off between price, performance and power dissipation [43]. The definition of an ultra-low power system is such that its price and performance is traded-off in favour of minimal power dissipation. With an idea of what an ultra-low power system is, the question arises as to how such a system is designed.

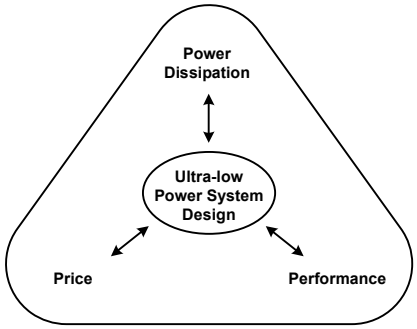

Figure 1.5: Design trade-offs associated with ultra-low power system design.

The guiding principles of ultra-low power system design impact all levels of design abstraction [16]. The minimisation of a system's power dissipation starts at the very lowest hardware representation (i.e. logic gate) and extends up to the highest level of software representation (i.e. software architecture). At each level of abstraction of a given system, one must evaluate intricate multi-objective optimisations so to determine the operation point that achieves power minimisation given the target application. The pay-off, or so called return on investment, will typically be greater the higher the level of abstraction is [16]. For example, optimising a software finite state machine may bring several orders of magnitude power improvement over a single gate optimisation in a microprocessor executing the software.

In the context of the prototype wireless hierarchical sensor network developed in this thesis, the power dissipation of the cluster node is the primary focus so to support extended operation lifetime. Therefore, the cost and performance (or feature set) of the cluster node is traded-off accordingly. The cost of the system is aligned towards commercial off-the-shelf (COTS) components in order to exercise rapid prototyping, whereas exuberant performance requirements (or features) have been relaxed in favour of minimising power dissipation. Starting with the most power efficient COTS components, a systematic bottom-up approach of minimising power, time and effort [16] is applied throughout the cluster node and cluster-head hardware and software design process.

1.4 Related Work

There are several commercial miniature temperature data loggers, including the iButton [84], the UTL-3 [70] and the Onset HOBO Pro [91], which have been successfully applied in Geoscientific research in [55], [59] and [44] respectively. Each of these devices are stand-alone temperature loggers with no support for wireless connectivity. The Onset HOBOnode [92] is a commercially available wireless temperature sensor, but its limited operation lifetime and lack of on-board memory make it unsuitable for scalable integration into a wireless sensor network.

The wireless hierarchical sensor network architecture exhibits similarities to disjoint clustering of stationary sensor nodes, with data aggregation at the cluster-head. However, since the size of each cluster is fixed and known based entirely on the target sensing application requirements, the cluster-head and all cluster nodes are stationary, and the cluster-head does not perform any topology control, there is no basis for incorporating dynamic clustering algorithms as surveyed in [28].

Wireless hierarchical sensing is very different from the functional decomposition exhibited in tiered sensor network architectures [27], as the cluster nodes

are considered to have fixed sensing, processing and communication capabilities. Furthermore, the spatial diversity achieved by the cluster nodes has no resemblance to the hierarchical sensing detailed in [53], as each cluster node supports a single-channel sensor which are aggregated using an energy efficient wake-up radio architecture.

The concept of a wake-up radio is certainly not new; it was first introduced in the literature [37] as early as 2001. Since then, there has been moderate growth in the field with contributions in both hardware and software-based proposals. A selection of notable contributions include the design and development of a custom wake-up radio ASIC based on FSK modulation having a power dissipation of 126μW [57], with PWM modulation dissipating 20μW [54], with OOK modulation dissipating 12.5μW [41], 10μW [60], 2.4μW [50], and with GOOK modulation dissipating 0.27μW [58]. Prototype systems combining active RF front-ends with microcontroller-based wake-up decoders based on OOK modulation having a maximum power dissipation up to 819μW [67], and 12.52μW [30]. Prototype systems with high power dissipation but with a low receiver sensitivity down to -122dBm [63]. Prototype systems [56] and [45] interfacing a microcontroller to COTS low-frequency wake-up receivers [75] and [76] respectively.

A wake-up radio architecture presented in [45] is used as a wake-up radio reference design in this thesis work. The paper demonstrates the coupling of a commercial low-frequency wake-up receiver to a passive OOK demodulator circuit. The paper also details how a commercial 868MHz ISM band radio module can be controlled to generate a wake-up sequence capable of triggering the wake-up receiver. The key differentiation between the wake-up radio implementation in [45] and this thesis is that the carrier frequency has been shifted down to 434MHz (for reasons detailed in Section 2.1.3), which required a complete re-design of the RF front-end circuitry.

According to the best knowledge of the author, this work represents the first attempt at developing a functional prototype of an ultra-low power wireless sensor node with an integrated wake-up radio based on COTS components, which provides periodic sensing, and that is capable of fully integrating into an existing real-world wireless sensor network infrastructure.

Chapter 2

System Requirements Analysis

2.1 Communication

2.1.1 Architecture

Wireless hierarchical sensor networking consists of two tiers of wireless connectivity. The upper-tier network interconnecting cluster-heads is a multi-hop network where media access control, mobility management, routing, flow control and data forwarding services are supported for end-to-end data delivery. This tier is implemented using energy efficient sensor network protocol stack (e.g. Dozer [35]). The lower-tier network between the cluster-head and a set of cluster nodes consists of point-to-point links in a star topology, and must support data delivery from each cluster node toward the cluster-head. The following section investigates energy efficient schemes for this lower-tier network.

At the core of the wireless hierarchical sensor network is a data delivery mechanism. Each cluster node accumulates measurement data (e.g. ground surface temperature measurements) in which the cluster-head is to aggregate together and forward over the upper-tier multi-hop network. A prerequisite for data aggregation is that the cluster-head is listening to the wireless channel when each cluster node is sending measurement data. This requires a rendezvous scheme, of which there are three types [53], namely pure synchronous, pseudo-asynchronous, and pure asynchronous. Figure 2.1 illustrates an abstract representation of these rendezvous schemes for data delivery between a set of cluster nodes and the cluster-head.

The pure synchronous scheme represents the communication architecture adopted in the majority of modern wireless sensor network protocol stacks, and thus is the basis of the aforementioned upper-tier multi-hop network. If this scheme is also applied to the lower-tier, data is transferred between the cluster node and cluster-head with the assistance of a bi-directional control channel. Typically this bi-directional control channel is used to configure a time synchronisation primitive, which is then used the schedule the data transfer. Maintaining

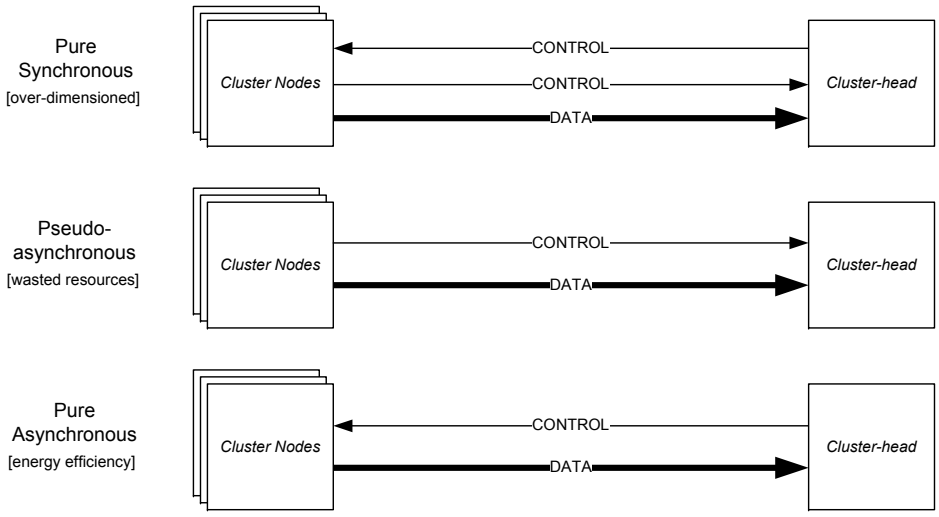

Figure 2.1: Abstract representation of rendezvous schemes for data delivery between a set of cluster nodes and a cluster-head.

the time synchronisation primitives on the cluster-head and cluster node requires considerable energy. Furthermore, since only minimal services are needed between the cluster nodes and the cluster-head (i.e. single-hop data transfer), the pure synchronous scheme is considered over-dimensioned for the lower-tier network.

In the pseudo-asynchronous scheme, each cluster node simply transmits as soon as it has data ready to send. The cluster node emits a preamble prior to sending the data to the cluster-head. The cluster-head has no way of knowing when data is to be sent from the cluster, which leads to wasted power resources though idle listening [9]. This scheme also has the potential for collisions during transmission as the nodes within the cluster are not time synchronised.

In the pure asynchronous scheme, the cluster-head is responsible for requesting data delivery through the transmission of a wake-up sequence. Each cluster node listens for the wake-up sequence using an always-on ultra-low power receiver. Upon reception of the wake-up sequence, each cluster node constructs a packet and invokes a multi-user access policy (i.e. TDMA) before the data is transmitted to the cluster-head. The cluster-head does not waste energy overhearing, there are no collisions between cluster nodes due to implicit time synchronisation, the cluster-head reduces idle listening, and there is minimal protocol overhead in its implementation. In summary, the pure asynchronous scheme is an energy efficient rendezvous scheme well suited for data delivery between a cluster and cluster-head, and thus forms the basis for the lower-tier network of the wireless hierarchical sensor network architecture.

2.1.2 Physical Channel Characterisation

Wireless sensor networks are predominantly implemented using far-field (i.e. radio frequency) communications, however there are research initiatives that have investigated the application of near-field (i.e. magnetic inductive coupling) communication techniques [29]. As detailed in [6], there are several physical limitations of near-field communications including limited range, polarisation effects, flux distortion and high power consumption, which collectively render near-field communications unsuitable for the prototype wireless hierarchical sensor network.

Radio frequency communication is therefore chosen as the means of wireless communication. The license-free spectrum in the Industrial, Scientific and Medical (ISM) radio bands are leveraged, in particular the 434MHz, 868MHz and 2400MHz frequency bands. These harmonised frequency bands are available throughout most of Europe and have their frequency bands, power levels, channel spacing and transmission duty-cycle regulated by CEPT/ERC [80] and ETSI [81, 82].

It is well known that adequate signal reception in a wireless communication system is achieved when the length of the antenna is proportional to the wavelength of the carrier frequency [12]. This means smaller antenna form factors can be adopted if a higher carrier frequency is used. However, due to the outdoor nature of the ground surface temperature use case, the length of the antenna is not a critical design constraint, thus all three aforementioned ISM frequency bands are feasible for the prototype wireless hierarchical sensor network.

The wireless communication of the hierarchical sensor network consists of two logical channels, namely a wake-up channel and a data channel. The two logical channels are mapped onto the same underlying physical channel, with differing symbol rates and channel encodings. Given a typical deployment scenario for a wireless hierarchical sensor network, the communication channel between the cluster-head and the cluster nodes is characterised as follows:

- *Asymmetric*: A control indication (i.e. the wake-up sequence) is initiated from the cluster-head, while measurement data is sent from the cluster node. There is no per-packet flow control.
- *Low Range*: The communication range of the cluster nodes is less than the cluster-head over the upper-tier wireless sensor network.
- *Low Data Rate*: The transmission data rate is low (i.e. O(kbps)).
- *Single-hop*: Communication is only supported on a single-hop (or point-to-point link) between cluster-head and cluster node.
- *Stationary*: The cluster-head and the cluster nodes are considered to be deployed in a fixed and stationary location.

2.1.3 Link Budget Analysis

A link budget analysis is used to quantify the functional requirements of a wireless communication system given the characteristics of the underlying physical channel. The modulation scheme, path loss model and target receiver sensitivity all impact the link budget analysis, and will each be treated in the following subsections.

2.1.3.1 Modulation Scheme

The selection of a communication system's modulation scheme is a trade-off between energy per bit and the bandwidth efficiency [13], based on Shannon's channel capacity theorem [62]. A modulation scheme with a lower energy per bit can deliver more data for a fixed amount of energy, whereas a modulation scheme with greater bandwidth efficiency can deliver a higher transmission rate for fixed amount of bandwidth.

The most common modulation schemes for low power communication systems are On-Off Keying (OOK), Frequency Shift Keying (FSK), and Phase Shift Keying (PSK) [13]. Figure 2.2 illustrates the bit error probability versus the energy per bit, or otherwise termed signal-to-noise ratio (SNR) per bit, for these modulation schemes. An extensive analysis of the these modulation schemes is found in [14].

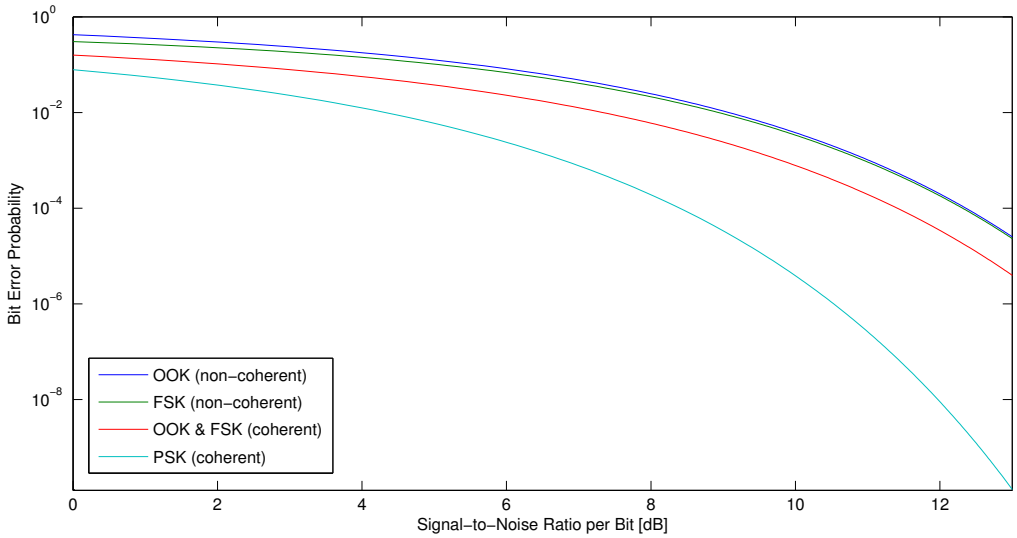

Figure 2.2: Error performance of OOK, FSK, and PSK modulation schemes, using coherent and non-coherent receiver structures.

Given the ultra-low power design constraints of the wireless hierarchical sensor network, it is justifiable to sacrifice some bandwidth efficiency in return for a reduced minimum achievable energy per bit. However, one must also consider the energy cost of the transmitter and receiver implementation. For example, PSK modulation features a reduced SNR per bit compared to FSK and OOK, as illustrated in Figure 2.2. However, the energy cost of a coherent PSK demodulator is relatively high due to the active components (i.e. local oscillator and phase-locked loop) needed to recover the phase information from the modulated symbols. In contrast, FSK and OOK receivers can be implemented with a non-coherent receiver structure, which does not rely on such power-hungry active components, therefore significantly reducing the energy cost of reception. For these reasons, PSK modulation is not considered as a viable modulation scheme, while the coherent implementations of FSK and OOK are not considered energy efficient receiver structures for the wireless hierarchical sensing system design.

The non-coherent FSK and OOK modulation schemes represent a compromise between bandwidth efficiency, minimum achievable SNR per bit, and energy cost of implementation. However, an analysis of their respective constellation diagrams as illustrated in Figure 2.3 indicates further energy savings.

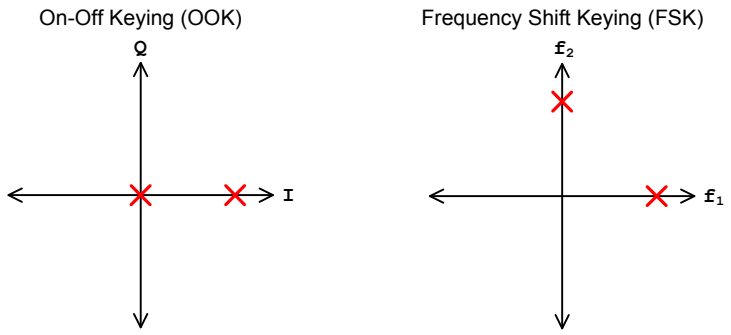

Figure 2.3: Constellation of (left) OOK in phase space, and (right) FSK in frequency space.

Assuming single bit per symbol mapping with equal symbol probabilities, OOK modulation only transmits a carrier frequency on half of the symbols on average compared to FSK modulation. Under these assumptions, theoretically one could expect a 50% reduction in energy consumption using the OOK modulation scheme compared to FSK. Therefore, the non-coherent OOK modulation scheme is chosen for the wireless hierarchical sensor network system design.

2.1.3.2 Path Loss Model

The path loss due to radio wave propagation is a fundamental phenomenon impacting all modern communication systems. The path loss contributors include free-space loss, absorption losses, multi-path fading losses and other impairments [18]. The aim is to evaluate the total path loss experienced between the cluster-head and the cluster nodes given the ground surface temperature use case.

It is expected that the cluster nodes in the wireless hierarchical sensor network are deployed within close range of, and typically within line-of-sight to, the cluster-head in an outdoor open area. The Friis free-space loss model [18] (see Equation 2.1) is used to model the path loss over the line-of-sight propagation path. The path loss is dependent on the line-of-sight distance d in meters, the carrier frequency f in Hertz, the unit-less path loss exponent η, and the speed of light c in m/s. A loss exponent of $\eta = 2$ is selected for the Friis model based on the deployment of the network in an outdoor open area [18].

$$L_{path} = 20\log\left(\frac{c}{f}\right) - 20\log(4\pi) - 10\eta\log(d) \tag{2.1}$$

The fading characteristics observed between the cluster-head and the cluster nodes can be determined by considering the coherence bandwidth and coherence time of the underlying physical channel [22]. Since the wireless hierarchical sensor network is considered stationary (see Section 2.1.2), the Doppler spread experienced at the cluster-head and cluster nodes is negligible. This results in a long coherence time of the wake-up and data channels. A survey of channel measurements [33] indicates that the delay spread of the multi-path components in sensor networks may be in the range of 10ns to 100ns. As assumed in Section 2.1.2, low data rates on the order of kilobits per seconds are to be expected over the wake-up and data channels. This implies that the coherence bandwidth of the physical channel is several orders of magnitude higher than the symbol bandwidth carried over the wake-up and data channels. This brief analysis suggests that the fading between the cluster-head and cluster nodes is considered slow frequency-flat fading [22]. Therefore, based on a similar analysis performed in [18], the slow multi-path fading can be factored into the path loss model through a 2dB loss in the link budget.

It is likely that the wireless hierarchical sensor network will be subject to harsh environmental conditions such as rain, snow and ice. The RF propagation effects of rain tend to only impact high-frequency ranges (i.e. gigahertz and higher) and over long distances (i.e. kilometers and longer), making their impact on the link budget negligible. The attenuation due to snow is primarily attributed to moisture content [48], and therefore expected to be even less than the losses brought by rain. Studies [42] suggest the absorption effects of ice are quite low (i.e. less than 1dB at a depth of 100m at 450MHz), and the loss by reflection

at the ice-air boundary is also less than 1dB at the carrier frequency of interest. Taking a rather conservative approach, the combined effects of rain, snow and ice can be incorporated into the path loss model with a 2dB loss in the link budget.

2.1.3.3 Receiver Sensitivity

The receiver sensitivity is defined as the minimum signal strength at the receiver which supports successful communication. A previous research project [45], from which the wake-up radio design is adapted from, claims a receiver sensitivity down to −52dBm on the 868MHz ISM band. It is envisaged that a similar sensitivity level is realistic, and serves as the receiver sensitivity goal for the wake-up radio in this thesis. A survey of modern commercial radio modules (see Appendix B.2) suggests that data radios with a receiver sensitivity of −110dBm is a realistic expectation.

2.1.3.4 Link Margin

The link budget (see Appendix A for details) is evaluated using a link margin expressed in dB. The link margin is the difference between the signal level at the receiver input and the receiver sensitivity, as expressed in Equation 2.2. The link margin M_{link} is dependent on the carrier frequency f in Hertz and the communication range d in meters. Communication is feasible over the wake-up and data channels when the respective link margins are positive. The aim is to determine a communication range d over the three ISM frequency bands of interest $f \in \{434, 868, 2400 \text{ MHz}\}$ such that a positive link margin is achieved on both wake-up and data channels.

$$M_{link} = P_{tx} + G_{tx} - \sum L_{tx} + L_{path}(f, d) - \sum L_{channel} + G_{rx} - \sum L_{rx} - S_{rx} \quad (2.2)$$

Figure 2.4 illustrates the wake-up and data channel link margins evaluated for the three ISM frequency bands of interest. It is clear from the link margin that the wake-up channel is the limiting link in the communication system. This is to be expected considering the low receiver sensitivity target of the wake-up radio compared to the data channel (see Section 2.1.3.3).

The results illustrated in Figure 2.4 indicate that the wake-up channel can be supported on the 434MHz carrier frequency up to a range of approximately 30 meters, as determined by the +2.1dB link margin. Therefore the 434MHz ISM band is selected for the prototype hierarchical wireless sensor network, with an aim to achieve up to 30 meters communication range between the cluster-head and cluster nodes.

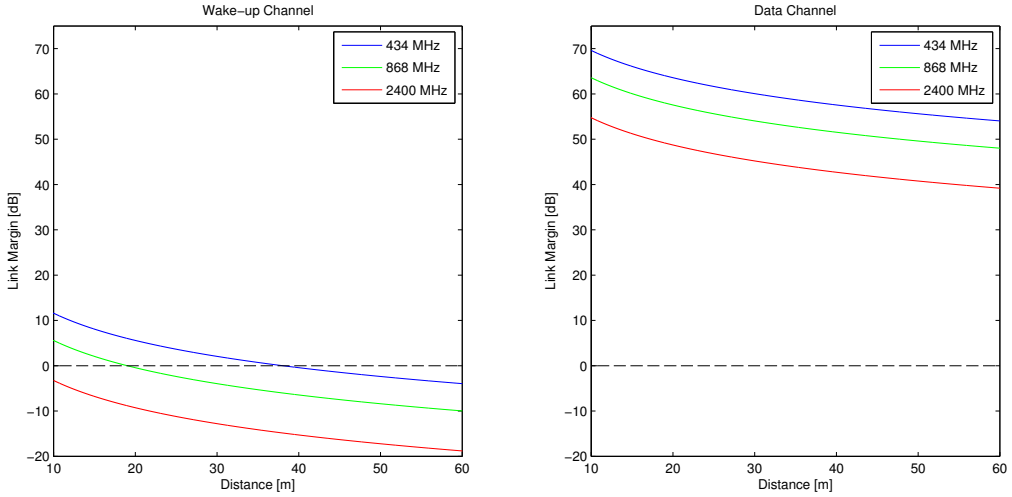

Figure 2.4: Link margin for (left) wake-up and (right) data channels.

2.2 Sensing & Processing

2.2.1 Temperature Sensor Specification

Due to the scientific nature of the ground surface temperature monitoring use case, the measurements collected by each cluster node need to be of sufficient quality. The quality relates to the temperature accuracy and resolution over the temperature range of interest. Based on the specifications of commercial temperature data loggers previously used in Geoscientific research [84, 70, 91], realistic temperature sensor specifications are summarised in Table 2.1.

Sensor Specification	**Value**
Measurement Range	-40 to $+80°C$
Accuracy	$\pm 0.2°C$
Resolution	$0.05°C$

Table 2.1: Sensor specifications for the ground surface temperature use case.

The temperature range and accuracy is a property of the selected temperature sensor. Temperature sensors can be partitioned into three main categories, namely electro-mechanical, electronic and resistive sensors [24]. The electro-mechanical sensors are typically large and cumbersome to integrate into small form factors whereas the electronic sensors are typically power inefficient and are costly to procure. Resistive sensors on the other hand are small, sensitive, cheap and can be powered efficiently, making them the ideal choice for integration into the cluster node.

The temperature resolution is a property of how the measurement sample is represented digitally. An Analog-to-Digital Converter (ADC) is used to digitise the analog signal representing the measured temperature. The temperature resolution determines the minimum number of bits the ADC must support per measurement sample. Assuming a linear temperature response (which is unlikely in practice, but helps guide the functional requirement analysis), the measurement resolution Q can be analytically evaluated by considering the full-scale ADC input voltage $V^+ - V^-$ and the number of bits per sample m, as defined in Equation 2.3.

$$Q = \frac{V^+ - V^-}{2^m - 1} \tag{2.3}$$

The positive ADC reference voltage V^+ is at most equal to the regulated supply voltage of the cluster node. The regulated voltage must be lower than the nominal battery voltage (e.g. 3.0V assuming a CR2477 battery [94]), and must be within the supply voltage range of all hardware components. A 2.5V regulated supply voltage is the lowest supply voltage that supports all cluster node hardware components.

Assuming the full range of the ADC (i.e. $V^- = 0$V, $V^+ = 2.5$V), Equation 2.3 can be solved in terms of m such that $Q \leq 0.05$°C. This implies the cluster node must incorporate at least an 11-bit ADC in order to support the sensing specifications of the given use case. Incorporating a design margin for temperature sensor non-linearity, a 12-bit ADC is preferred.

2.2.2 Microcontroller Considerations

The microcontroller must perform several tasks at both absolute time (e.g. periodic measurements) and relative time (e.g. TDMA-slot alignment) resolution. An internal or external real-time clock with calendar support is needed for absolute time maintenance with resolution of at least one minute. One or more internal 16-bit timers are needed to support relative time triggering with at least one millisecond resolution.

The choice of the microcontroller architecture (i.e. 8, 16, or 32-bit) is not critical to the overall operation of the cluster node. However, given the temperature resolution is greater than 8-bits, and that the likelihood of utilising at least one 16-bit compatible SPI interface is high, there are implementation advantages in selecting a 16-bit or higher architecture.

The sleep performance of the microcontroller is vital to the cluster nodes overall power dissipation. The ability to wake-up from sleep mode using a Real-Time Clock (RTC) alarm or an external interrupt will not only help minimising the overall power dissipation, but will also greatly simplify the software application logic.

Finally, the microcontroller must support the desired external peripherals, with at least one Universal Asynchronous Receiver Transmitter (UART) interface and additional General Purpose Input Output (GPIO) pins for on-board hardware and software debugging.

2.3 Data Management

2.3.1 Data Aggregation & Multi-user Access

The wireless hierarchical sensor network supports delivery of sensor data from many cluster nodes to a single cluster-head. There are two main schemes for facilitating this data collection, namely using a push-based or a pull-based data collection mechanism [15]. There are two compelling reasons for selecting a pull-based data collection mechanism, which will be described in the following paragraphs.

The underlying physical process being measured (e.g. ground surface temperature) changes slowly. This means the time between successive temperature measurements needs not be faster than several minutes. These two factors imply a low data rate with relaxed data latency constraints. According to [51], efficient data retrieval can be achieved using a pull-based data collection mechanism due to the implied low data latency constraints.

If a push-based data collection mechanism is considered instead, all cluster nodes would need to be time synchronised so to avoid collisions during transmission. Achieving this time synchronisation induces protocol overhead, which translates into higher energy consumption. However, assuming a pull-based data collection mechanism, all the cluster nodes are implicitly synchronised through the wake-up sequence. As the communication range is low, the propagation time of the incident electromagnetic wave between cluster-head and cluster node will be on the order of a few microcontroller clock cycles. This results in the wake-up interrupt triggering the cluster node microcontroller at essentially the same time for all nodes in the cluster. Thus, the pull-based mechanism is the superior choice.

In order to differentiate multiple data streams at the cluster-head, each cluster node must have a unique identifier. In a realistic deployment scenario, it is unlikely that more than 256 cluster nodes would be placed in the vicinity of a single cluster-head. Therefore the resolution of the unique cluster node identification is limited to 8-bits.

A multi-user access policy must be enforced on each cluster node so to prevent collisions during data transmission.

2.3.2 Data Redundancy & Validity

Data redundancy and validity are very important features in any data analysis application. Due to the nature of wireless communications and the long-term stability of embedded systems, seemingly unexplainable circumstances may arise where measurement data is never requested, transmitted nor received from a given cluster node. It is therefore imperative that all data measurements are locally stored in non-volatile memory. This enables the measurement data to be manually recovered for data redundancy and validity purposes.

2.4 Operation & Maintenance

2.4.1 Operation Lifetime

In order to support long-term environmental monitoring, it is the aim that the cluster nodes have an operation lifetime of several years on a single battery charge. Commercial temperature data loggers [84, 70, 91] specify field durations in the order of $3 - 5$ years, at sampling intervals of $1 - 4$ hours. Through consultation with Geoscience researchers, a cluster node operation lifetime of 3 years at a 1 hour measurement interval serves as a realistic functional requirement. Flexible deployment use cases are envisaged through the support of measurement intervals between 10 minutes and 4 hours.

2.4.2 Deployment Considerations

The cluster node requires local configuration prior to deployment. At the very least, the cluster node must have the current time and the measurement interval configured. It is also desirable to read-out measurement data after the manual collection of devices from the field. To simplify deployment, it is advantageous to support a human-readable configuration menu using a platform-independent interface.

The cluster-head must deliver received measurement data via a data interface. The type of data interface depends on how the cluster-head is integrated into the target application. A human-readable data interface is most suitable for the purposes of a prototype implementation.

The prototype of the wireless hierarchical sensor network must be constructed from COTS components. Furthermore, the temperature range of all electronic parts must be within the range of the temperature sensor (i.e. -40 and $+80°C$).

2.5 Summary of Functional Requirements

The requirements analysis presented in this chapter defines the functional requirements for the prototype wireless hierarchical sensor network developed in this thesis. However, with the exception of the sensor specifications (see Section 2.2.1), the functional requirements are independent on the application use case. This ensures that the prototype wireless hierarchical sensor network can be adapted into a multitude of application domains.

The functional requirements of the prototype ultra-low power wireless hierarchical sensor network, assuming the ground surface temperature use case, are as follows:

Communication:

- Wireless reception of measurement data using a pull-based data collection mechanism.
- A point-to-point communication range of 30 meters on the 434MHz ISM frequency band in outdoor conditions.

Sensing & Processing:

- Temperature measurement range between -40 and $+80$°C.
- Temperature accuracy of at least ± 0.2°C over the majority of the temperature range.
- Temperature resolution of at least 0.05°C over the majority of the temperature range.
- Configurable periodic measurement sampling rate between 10 minutes and 4 hours.
- Each measurement has an associated acquisition time with at least one minute temporal resolution.

Data Management:

- All measurement data stored in non-volatile memory.
- Each cluster node has a unique system identification number.
- A maximum of 256 cluster nodes supported in a single cluster.
- Ability to set system time and measurement interval, and download all stored measurement data using a human-readable and platform-independent configuration interface.
- Cluster-head provides a generic data interface for delivery of measurement data.

Operation & Maintenance:

- All system components rated for temperatures between -40 and $+80$°C.
- Support for basic on-board hardware and software debugging.
- Cluster node powered by a single coin cell battery.

Design Constraints:

- System designed using COTS components.
- Cluster node and data lifetime of at least 3 years at a 1 hour sample interval.

Chapter 3

System Design & Implementation

3.1 System Architecture

The overall system architecture of the cluster node and the cluster-head is illustrated in Figure 3.1. The following sections detail the design and implementation of the communication, sensing, processing and data management functions.

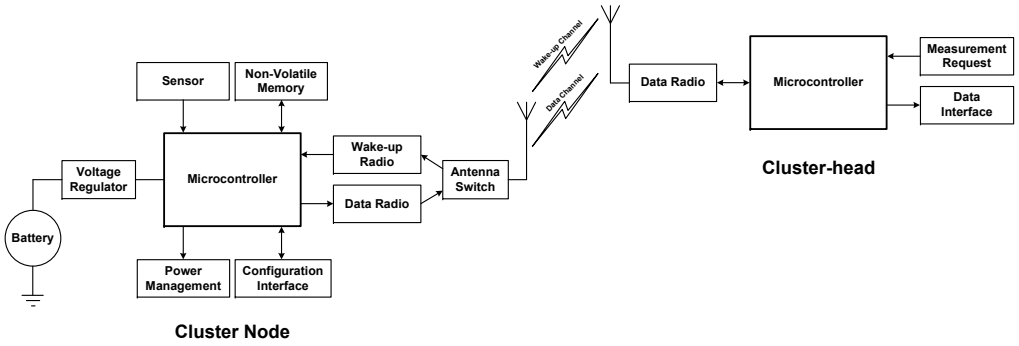

Figure 3.1: System architecture of cluster node and cluster-head.

3.2 Communication

The data flow of the communication between the cluster node and the cluster-head is illustrated in Figure 3.2. Communication is initiated from the cluster-head by requesting data from the cluster nodes within its vicinity. This is achieved by generating an appropriate wake-up sequence (see Section 3.2.4.1) and transmitting it using its data radio. The data radio transmits the wake-up sequence on the 434MHz carrier frequency using On-Off Keying (OOK) modulation.

The cluster node receives the wake-up sequence through its antenna and antenna switch (see Section 3.2.1) and passes it through to the wake-up radio. The wake-up radio consists of a custom On-Off Keying demodulator circuit (see Section 3.2.2) coupled with a commercially available 125kHz low-frequency wake-up receiver (see Section 3.2.3). The design of the wake-up radio is an adaption from [45]. The OOK demodulator recovers the wake-up sequence and passes it to the wake-up receiver for processing. Upon successful detection of the wake-up sequence, an interrupt is passed to the microcontroller awaking it from a low power sleep mode. The microcontroller constructs a data packet containing previously stored sensor data and transmits the packet using the data radio (see Section 3.2.4.2). The cluster-head receives the packet using its data radio and presents it to the data interface (see Section 3.4.4) for target application processing.

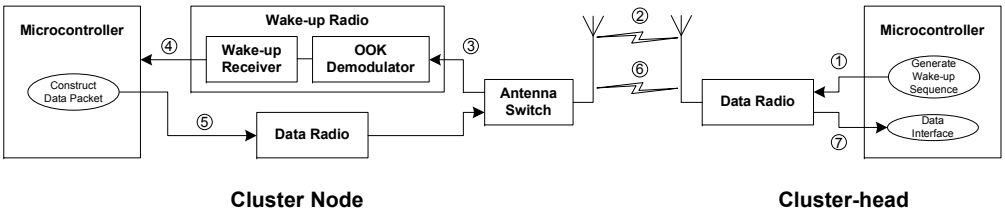

Figure 3.2: Communication data flow between cluster node and cluster-head.

3.2.1 Antenna & Antenna Switch

The wake-up and data channels are both centered onto the 434MHz carrier frequency and use a narrow-band modulation scheme. A single antenna can therefore be shared by the cluster node for both channels through an antenna switch. The ADG918 antenna switch from Analog Devices [73] is used to control the cluster node's access to the wake-up and data channel. The ASG918 is an absorptive switch with 50Ω terminated inputs, and features a low quiescent current drain of 0.1μA. The antenna switch is operated using a digital control pin on the microcontroller.

Monopole, dipole and helical antennas tuned to the 434MHz ISM band with 50Ω impedance are all viable candidates to feed the wake-up and data radios. The bandwidth of the antenna will vary greatly depending on the manufacturer and the specific ground plane used with the antenna. Through experimentation, a 433MHz quarter-wavelength monopole whip antenna from AntennaFactor [74] is chosen for the design.

3.2.2 On-Off Keying Demodulator

The purpose of the OOK demodulator is to recover the original 125kHz square-wave signal encoding the wake-up sequence from an OOK modulated 434MHz carrier frequency.

There are several OOK demodulator receiver architectures in the literature, including Amplitude Modulation (AM) envelope detectors, direct-conversion receivers, and superheterodyne receivers [3]. Given the ultra-low power requirements of the cluster node, the AM envelope detector receiver structure is selected due to its passive RF front-end circuitry. The OOK demodulator structure is illustrated in Figure 3.3. As previously stated, the OOK demodulator is an adaptation of the design presented in [45].

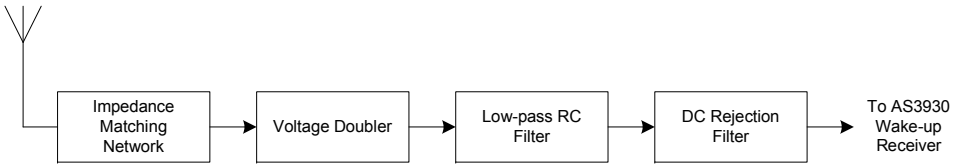

Figure 3.3: OOK demodulator structure.

An antenna provides the 434MHz RF input into the OOK demodulator. An impedance matching network is used to match the standard 50Ω antenna impedance to the impedance of the remaining RF circuit stages. The received signal is then passed into a voltage doubler circuit, which rectifies the sinusoidal input signal producing a voltage proportional to twice the peak of the input signal. The rectified signal is then filtered using a low-pass RC filter, followed by the removal of any DC offset using a high pass filter. The output signal is an approximate representation of the 125kHz square-wave signal encoding the wake-up sequence. This signal is then passed to the wake-up receiver for detection.

The OOK demodulator is built from common passive and semi-conductor components (i.e. resistors, inductors, capacitors and diodes), and since the circuit is operating at high frequencies, the component selection and detailed circuit design is complex. The final circuit design is realised through a combination of circuit analysis and synthesis techniques. The flow diagram in Figure 3.4 illustrates the lengthy design process and the analysis and synthesis tools used at each stage. Contrary to intuition, the design process of the OOK demodulator starts at the output and works backwards through the RF chain toward the input. The following subsections will detail each step in the design process in the order presented in Figure 3.4.

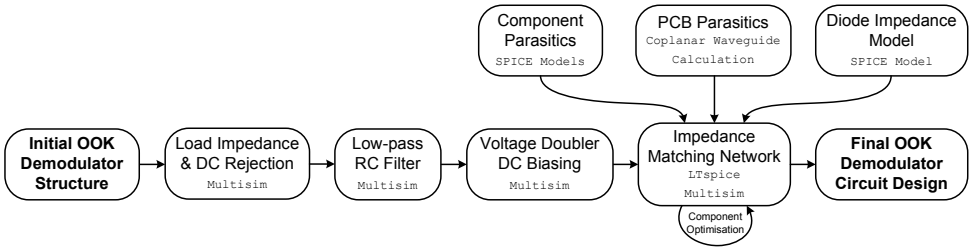

Figure 3.4: Design flow of the OOK demodulator circuit.

3.2.2.1 Load Impedance & DC Rejection Filter

The intended output of the OOK demodulator is the recovered 125kHz square-wave signal encoding the wake-up sequence. This signal is then passed to the AS3930 wake-up receiver for further processing. The AS3930 has a very high input impedance of 2MΩ [76], which is modelled as a purely resistive load impedance throughout subsequent circuit analysis.

The presence of a large DC offset will adversely affect the sensitivity of the AS3930 receiver, and therefore must be removed using a series capacitor, which forms a first-order high pass filter with the 2MΩ load impedance. The value of the DC rejection capacitor must be low enough so as not to discharge through the adjoining low-pass RC filter, and must be high enough to pass through the 125kHz signal of interest. Multisim simulations indicate that a DC rejection capacitor of 220pF is a suitable choice.

3.2.2.2 Low-pass RC Filter

It is the intention of the OOK demodulator to generate an approximate 125kHz square-wave signal corresponding to the input on and off-symbol periods of the OOK modulated 434MHz carrier frequency. This can be generated passively using a parallel capacitor and resistor, which together form a first-order low-pass filter. During the on-symbol period, the rectified output of the voltage doubler charges the capacitor C, and is then discharged through resistor R during the off-symbol period.

The value of capacitor C must be chosen such that it is greater than the parasitic capacitance of the diodes used in the voltage doubler circuit (see Section 3.2.2.3 for further details). Based on the reference design in [45] and through Multisim simulation, the value of $C = 18$pF is chosen, which is approximately an order of magnitude larger than the parasitic capacitance of the voltage doubler diodes.

In order to generate an approximate 125kHz square-wave, the charge time of capacitor C must be approximately equal to half a period of the 125kHz waveform. Equation 3.1 defines the voltage of a capacitor V_c as a function of time t, where $\tau = RC$ is termed the time constant [8].

$$V_c = V_s \left(1 - \exp\left(\frac{-t}{\tau}\right)\right) \tag{3.1}$$

The value of R and C is chosen such that the voltage on the capacitor V_c has reached 99% of the supply voltage V_o by the time equal to half a period of the ideal 125kHz square-wave signal (i.e. 4μs). Using $C = 18$pF and solving Equation 3.1 for R, adjusting to standard component values and finally verifying the circuit in Multisim, a value of $R = 51\text{k}\Omega$ is selected.

3.2.2.3 Voltage Doubler

The voltage doubler is a type of rectification circuit which has the advantageous property of producing a direct voltage equal to the peak-to-peak value of a given input sinusoidal voltage [2]. The voltage doubler circuit used in the OOK demodulator is illustrated in Figure 3.5.

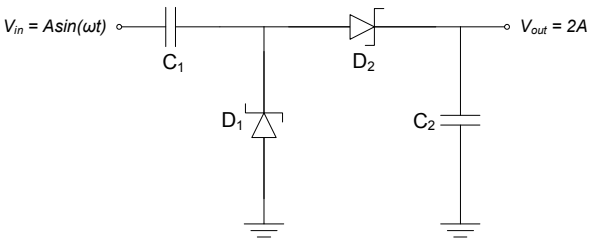

Figure 3.5: Voltage doubler circuit used in the OOK demodulator.

The voltage doubler circuit consists of a clamping diode D_1 and its charging capacitor C_1, in series with a detection diode D_2 and its charging capacitor C_2. The operation of the circuit is as follows; given a sinusoidal input signal, the capacitor C_1 will charge up to the positive peak voltage of the input sinusoidal. The voltage across the clamping diode D_1 will then be equal to the sinusoidal amplitude plus the charge stored in C_1, effectively "clamping" the sinusoid above ground. The detection diode D_2 then acts as a half-wave peak rectifier and charges the output capacitor C_2, producing a DC voltage (potentially with a small ripple) with double the peak-to-peak voltage of the input signal.

The clamping capacitor C_1 must be chosen such that the charging time of the capacitor is close to the period of the 125kHz signal of interest, and that

the capacitor has a low impedance at the 434MHz carrier frequency. A value of $C_1 = 39$pF is chosen based on the reference circuit in [45] and verified through Multisim simulation. The charging capacitor C_2 is equal to 18pF, as it is the capacitor associated with the RC low-pass filter detailed earlier in Section 3.2.2.2.

The primary function of a diode is to enable the forward flow of current across its junction depending on the applied forward voltage, as characterised by its I-V curve. Schottky diodes are a special type of diode consisting of a metal-semiconductor junction, which supports a low junction barrier voltage over a wide frequency range [25]. These properties make Schottky diodes favourable for several types of RF circuits, in particular the voltage doubler circuit used in the OOK demodulator.

The HSMS-2822 RF Schottky diode from Avago Technologies [77] are selected for diodes D_1 and D_2 in the voltage doubler circuit. The selection of the HSMS-2822 diode is based on previous research [45], the generality of the diodes specifications, and the availability of two series diodes in a single surface mount package. The I-V curve of the HSMS-2822 diode is illustrated in Figure 3.6.

Due to the physical properties of the materials used to construct semiconductor devices, each diode in the HSMS-2822 package exhibits a series resistance R_s, a junction resistance R_j, a junction capacitance C_j, a package capacitance C_p, and several leadframe L_l and bondwire L_b inductances. These parasitic properties are incorporated into a linear equivalent circuit model [79], as illustrated in Figure 3.6(right).

The HSMS-2822 diode has a very low, albeit non-zero, forward barrier voltage. This non-ideal behaviour adversely effects the operation of the voltage doubler circuit at low input signal levels. If the peak forward voltage from an input signal is not large enough to switch on the clamping diode D_1, the detection diode D_2 will rectify a very small sinusoidal, which will appear as an even smaller sinusoid at the output of the low-pass RC filter. This will significantly impact the performance of the OOK demodulator, and indeed the entire wake-up radio. To ensure the clamping diode D_1 is turned on (i.e. D_1 is forward biased) at low input signal levels, a DC bias must be applied.

The operation point of the clamping diode, that is the minimum input signal such that the clamping diode is forward biased, determines the voltage of the DC bias. As it is the aim to achieve a signal sensitivity down to -52dBm (equivalent to an ideal sinusoid with approx. 1mV amplitude), this equates to approximately 1mV / 50Ω = 20μA forward current at the input of the voltage doubler. This is a very optimistic approximation and seems unlikely in practice due to signal losses. However, this rough calculation gives an order-of-magnitude estimate for the operation point. Interpolating this operation point on the I-V curve in Figure 3.6 indicates that an approximate 200mV forward bias voltage is required across the clamping diode.

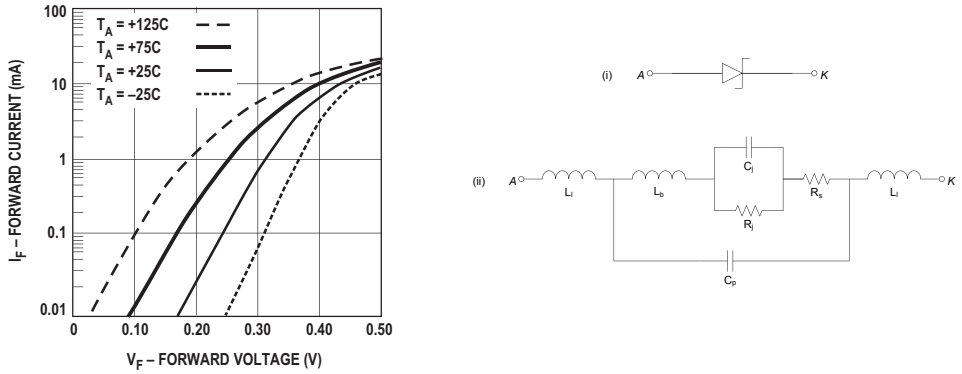

Figure 3.6: (left) The I-V curve for the HSMS-2822 Schottky diode, taken from [77]; (right)(i) the circuit symbol, and (ii) the linear equivalent circuit model for the HSMS-2822 Schottky diodes, adapted from [79].

The bias circuit is constructed from a resistive network coupled with a shunt capacitor to ground. The resistive network is designed such that it provides 200mV DC bias to the anode of the clamping diode D_1, while also acting as a high impedance at the 434MHz carrier frequency. The shunt capacitor provides a low impedance path to ground for the carrier frequency, ensuring the AC operation of the clamping diode is preserved. The diode bias circuit is illustrated in Figure 3.11.

3.2.2.4 Impedance Matching Network

The purpose of the impedance matching network is to facilitate the maximal power transfer from the antenna toward the rest of the OOK demodulator circuitry. It is well known from fundamental transmission line theory [1] that the power transfer between a source Z_{source} and load Z_{load} impedance is maximal when the load is equal to the complex conjugate of the source impedance. In practice, the source and load impedance of a circuit are rarely conjugately matched. In fact, there are several practical examples where a source impedance needs to be transformed into a load impedance for other purposes [1] (e.g. maximise the gain of a filter). The conversion of impedance between a source and load over a narrow bandwidth is realised through the application of circuit transformations using passive components.

At high-frequencies (i.e. above a few MHz), the behaviour of passive components such as resistors, capacitors and inductors no longer behave in an ideal way, as they are assumed to at low-frequencies [17]. The true behaviour of these passive components at high-frequencies are affected by so called parasitic re-

sistances, capacitances and inductances. These component parasitics introduce frequency dependent behaviour, which significantly impacts overall circuit operation. Therefore, it is vital to include such parasitic component properties in the circuit design through appropriate SPICE component models.

The behaviour of signals flowing around a circuit changes when operating at high frequencies. Signals flow according to the path of least resistance at low-frequencies. However, signals flow according to the path of least impedance at high-frequencies [5]. This phenomenon greatly impacts the design of Printed Circuit Boards (PCB) for high-frequency applications. The physical properties of the tracks interconnecting components impact overall circuit behaviour, and must therefore also be incorporated into the circuit design using appropriate equivalent lumped element models.

The input to the matching network is a standard antenna feed, making the source impedance of the matching network $Z_{source} = 50\Omega$. The load impedance must be determined through simulation of the RF-equivalent circuit incorporating the voltage doubler, the low-pass RC filter and the diode bias circuits. Figure 3.7 illustrates the Multisim simulation used to evaluate the load impedance, taking into consideration all component and PCB parasitics. A single frequency AC simulation at 434MHz yields a load impedance of $Z_{load} = 12.4957 - j214.4730\Omega$.

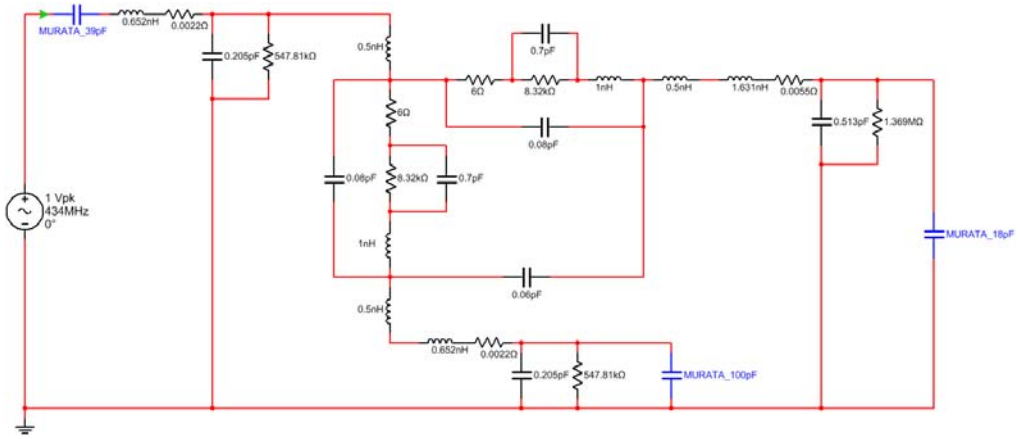

Figure 3.7: Multisim simulation used to calculate the load impedance of the RF-equivalent OOK demodulator circuit.

There are many passive impedance matching structures one can choose from to satisfy the impedance mismatch, including L, PI and T-networks [3]. An L-network coupled with the absorption technique [17] are used to match the resistive Z_{source} impedance to the resistive and reactive Z_{load} impedance found in the OOK demodulator circuit.

The impedance matching circuit is synthesised using a series of circuit transformations, as illustrated in Figure 3.8 and described in the following. The circuit synthesis begins with representing the resistive part of Z_{load} as a series resistor R_s, and the reactive part of Z_{load} as a capacitor C_s (see Figure 3.8(ii)). Since inductors have a positive reactance, a series inductor L^* is inserted between source and load so to absorb the stray reactance introduced by the load capacitor C_s (see Figure 3.8(iii)). The value of L^* is calculated such that its reactance is equal to the reactance of capacitor C_s at 434MHz. The remaining load resistance R_s can be matched to the 50Ω source resistance using an L-network. Since R_s is less than the source impedance, an L-network with an upward transformation is needed [1]. An L-network with a shunt-capacitor C_1 in series with an inductor L_1 is selected (see Figure 3.8(iv)). The values of C_1 and L_1 are calculated using formulae listed in [17]. In a final step, the two series inductors L_1 and L^* are combined, yielding the impedance matching circuit in Figure 3.8(v).

The values of C_1 and L_1 need to be adjusted according to standard component values, and then further optimised based on component and PCB parasitics. The component parasitics are incorporated through SPICE models made available through the component manufacturer. The PCB parasitics are modelled as a lumped element equivalent circuit, with the exact circuit depending on the type of transmission line selected for the PCB.

Coplanar waveguides [20] are selected as the transmission line type for the OOK demodulator PCB, due to their double-layer ground plane and the ability to efficiently machine the tracks using miniature slot-drill cutters. A coplanar waveguide calculator [69] is used to calculate the lumped element equivalent circuit based on the physical properties of the PCB and the dimensions of the machined tracks.

Figure 3.9 depicts the circuit used to simulate the impedance matching circuit in LTspice, incorporating both component and PCB parasitics. The aim of the simulation is to optimise the component values of C_1 and L_1 so to maximise the magnitude of the frequency response at 434MHz, measured at the input to the voltage doubler circuit (i.e. across capacitor C_1). The component optimisation process is one of trial and error. However, experience has shown that suitable component values are typically close to the calculated values. Figure 3.10 illustrates the simulated frequency response across C_1 using the refined component values of $C_1 = 5.6$pF and $L_1 = 82$nH. Figure 3.11 illustrates the final OOK demodulator circuit used in the Mutlisim simulation.

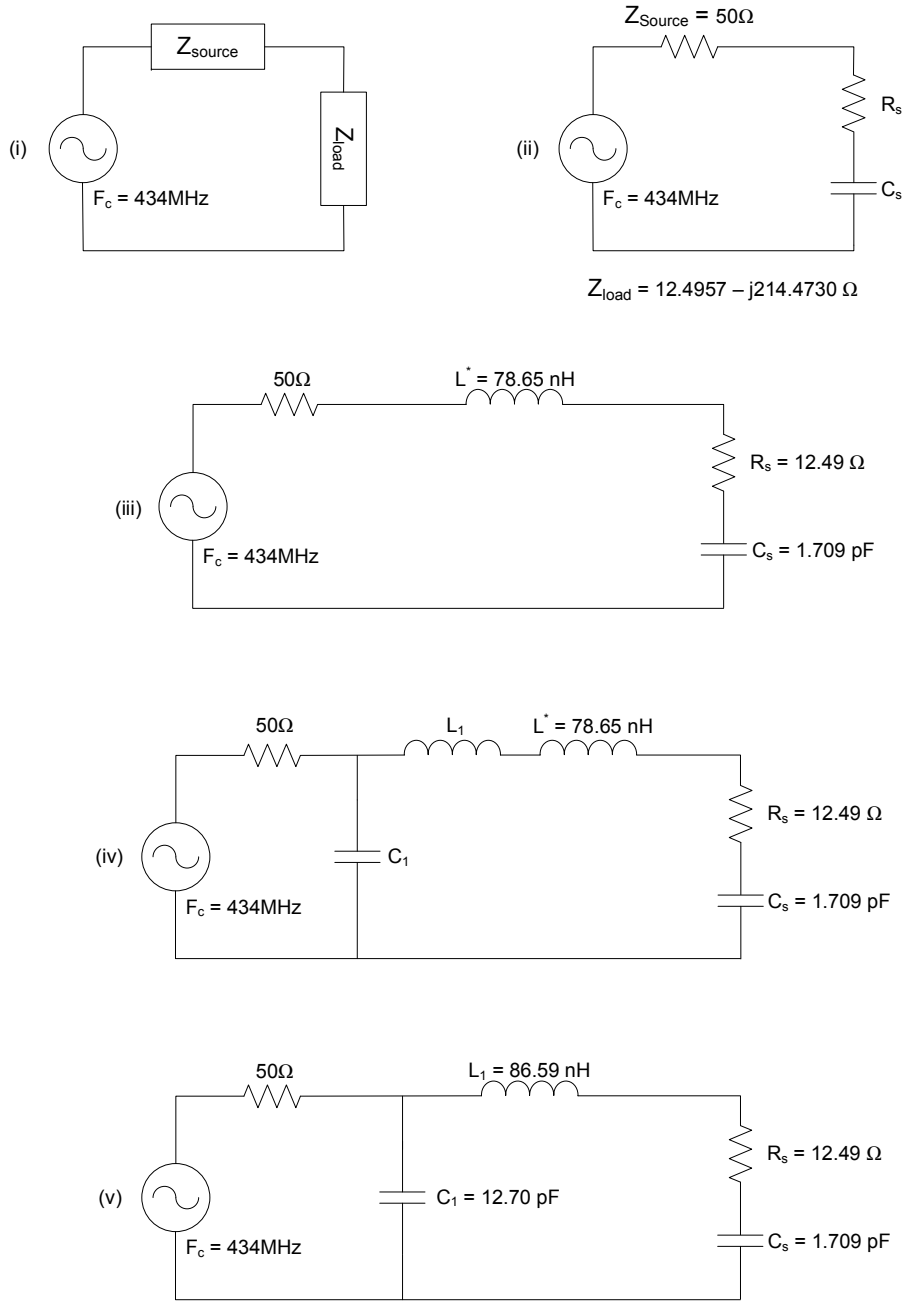

Figure 3.8: Synthesis of impedance matching circuit.

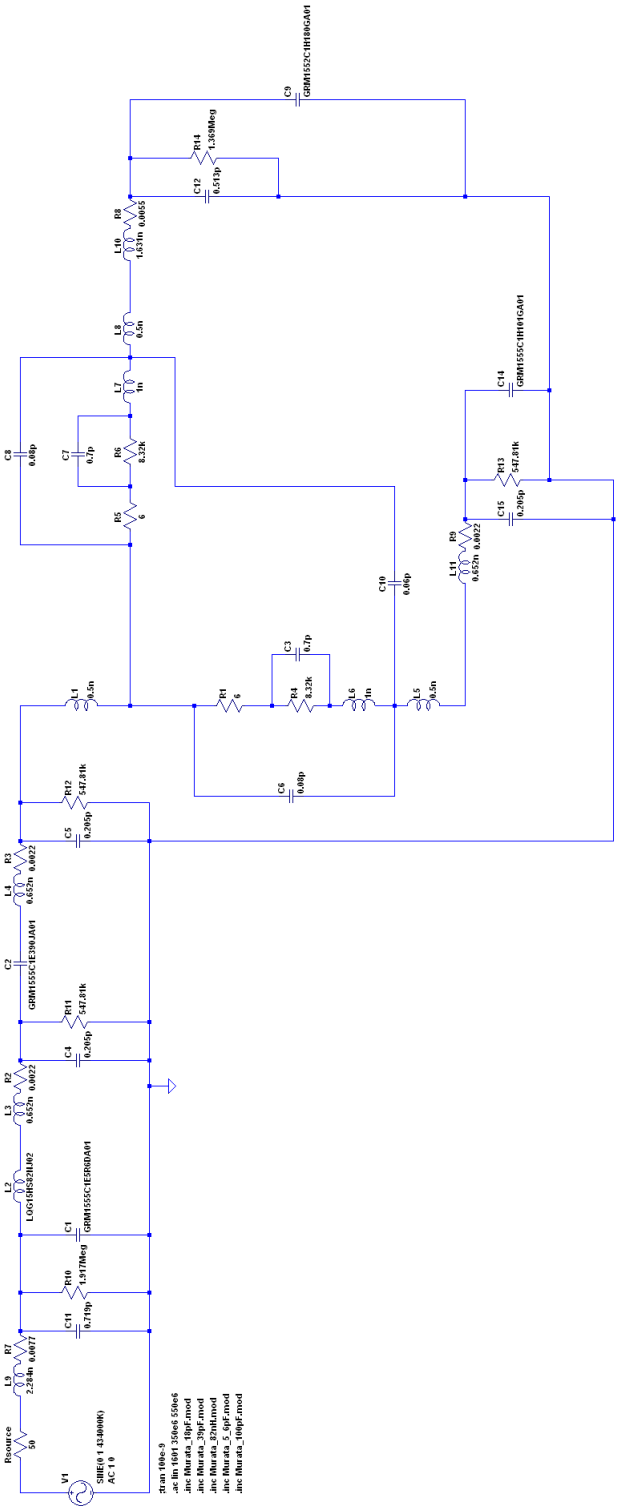

Figure 3.9: LTspice simulation of impedance matching circuit incorporating component and PCB parasitics.

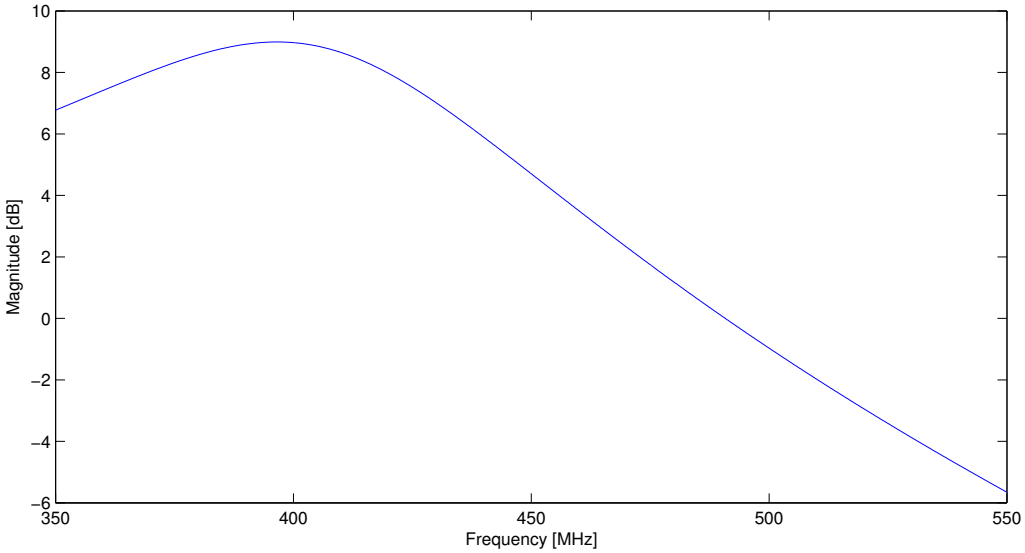

Figure 3.10: LTspice simulation of frequency response measured at the input to the voltage doubler circuit.

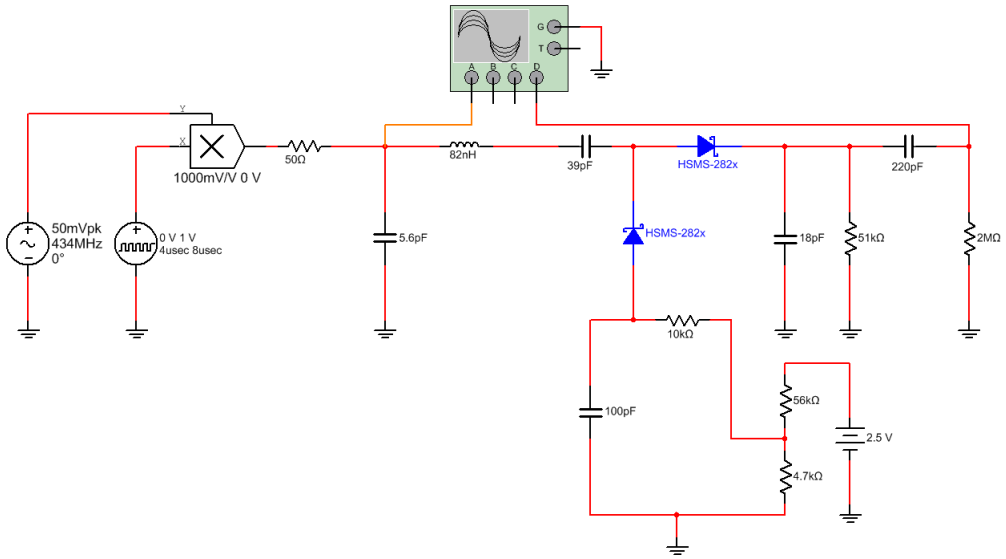

Figure 3.11: Multisim simulation of final OOK demodulator circuit.

3.2.2.5 Power-optimised On-Off Keying Demodulator

The DC bias current used by the OOK demodulator circuit is detrimental to the cluster node's power performance. The choice of Schottky diode used in the voltage double circuit determines the level of DC biasing needed to operate the OOK demodulator circuit. In particular, there is a range of commercial Schottky diodes, termed zero bias Schottky diodes, that do not require any DC biasing. Zero bias Schottky diodes are typically used for small signal RF detection circuits having low power input signals (i.e. less than -20dBm), making them well suited to the OOK demodulator circuit. Therefore, replacing the HSMS-2822 diodes for zero bias diodes would power-optimise the OOK demodulator circuit.

The HSMS-285C zero bias Schottky detector diodes [78] have been chosen to implement the power-optimised OOK demodulator, based on consultation with the author of [45]. The HSMS-285C zero bias diodes have a higher series resistance than the HSMS-2822 diodes, which significantly alters the load impedance of the impedance matching network. This necessitates a re-design of the impedance matching circuit, repeating the steps detailed in Section 3.2.2.4.

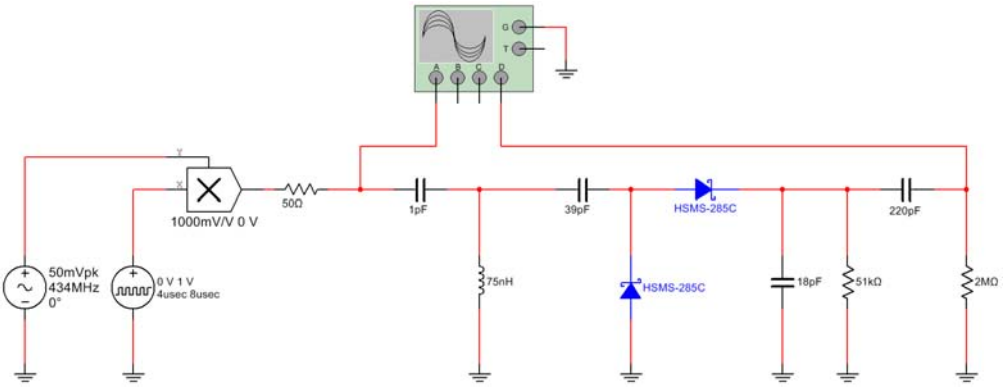

Figure 3.12: Multisim simulation of final power-optimised OOK demodulator circuit.

The total impedance of the voltage doubler circuit with the HSMS-285C diodes, the low pass RC filter and all component and PCB parasitics is evaluated using a Multisim simulation. A single frequency AC simulation at the carrier frequency of 434MHz yields a load impedance of $Z_{load} = 61.2588 - j506.9442\Omega$. Unlike the analog design in Section 3.2.2.4, the resistive component of Z_{load} is now larger than the purely resistive source impedance of $Z_{source} = 50\Omega$. This implies an L-network with a downward transformation [1] must be applied, consisting of a series capacitor $\bar{C}_1$ and a shunt inductor $\bar{L}_1$. The component values are calculated using known formulas in [17], adjusted to standard component values, and further refined due to component and PCB parasitics using an LT-

spice simulation. Figure 3.12 illustrates the final Multisim simulation with the refined component values for $\bar{C}_1 = 1\text{pF}$ and $\bar{L}_1 = 75\text{nH}$.

Figure 3.13 illustrates the simulated frequency response at the input to the voltage doubler circuit using the HSMS-2822 and HSMS-285C diodes. The simulation indicates that the new impedance matching circuit achieves a gain of more than 5dB at the 434MHz carrier frequency. This is very advantageous; not only is the new OOK demodulator circuit power-optimised by the removal of the DC bias, its receiver sensitivity is also increased through an improved impedance matching. These performance improvements are evaluated and quantified in Section 4.

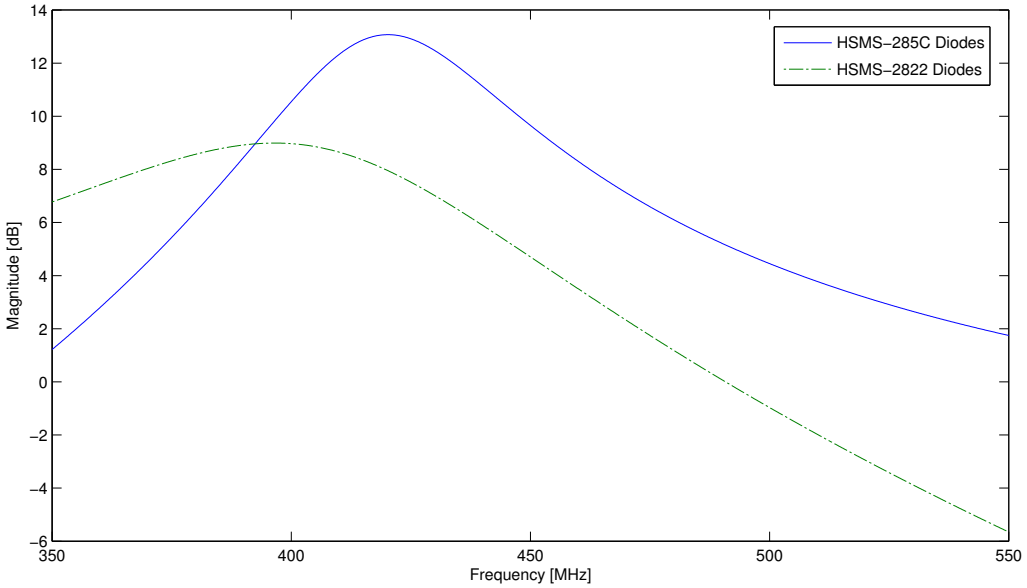

Figure 3.13: LTspice simulation of frequency response measured at the input to the voltage doubler circuit using zero biased and biased Schottky diodes.

3.2.3 Wake-up Receiver

The AS3930 single-channel 125kHz wake-up receiver from AustriaMicroSystems [76] forms the heart of the wake-up radio. A simplified block diagram of the device is illustrated in Figure 3.14. The AS3930 amplifies, detects and correlates a Manchester encoded 125kHz OOK input signal against a configurable 16-bit wake-up identifier. A wake-up interrupt signal is triggered upon the successful reception of the wake-up sequence.

The AS3930 has a configurable correlation symbol period making it possible to adapt the wake-up sequence to a high-frequency carrier (i.e. 434MHz) using

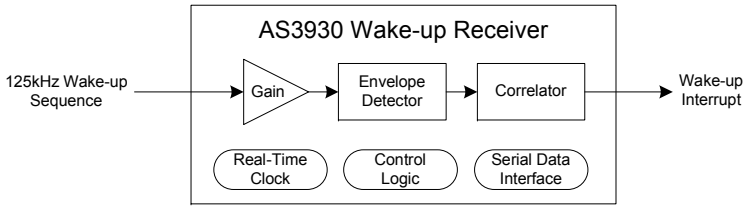

Figure 3.14: Simplified block diagram of AS3930 wake-up receiver.

OOK modulation. Furthermore, the device features a very low power consumption of 2.7μA [76], making it well-suited to ultra-low power wireless hierarchical sensing. The programmable 16-bit wake-up identifier provides a mechanism for supporting addressable wake-ups, however all cluster nodes in a wireless hierarchical sensor network are pre-configured with an identical 16-bit wake-up identifier.

The AS3930 is configured through a proprietary bi-directional Serial Data Interface (SDI), and signals the wake-up interrupt through an active high digital pin. The SDI interface is comparable to the standard SPI interface, with the exception of an active high chip select signal.

3.2.4 Data Radio

A data radio is incorporated into both the cluster node and the cluster-head, but with slightly different purposes. The cluster node uses the data radio to transmit measurement data after a wake-up sequence has been received (i.e. data radio in TX mode only), whereas the cluster-head transmits the wake-up sequence and receives measurement data using the data radio (i.e. data radio in TX and RX modes). Through an extensive design space exploration (see Appendix B for details), the CC115L receiver [97] and the CC110L transceiver [96] radio modules from Texas Instruments are selected as the data radio for the cluster node and cluster-head respectively.

The TI CC11xL are cost effective sub-1GHz radio modules supporting the 315/433/868/915 ISM bands with a wide range of data rates and modulation schemes. The radio modules are interfaced through a standard SPI interface for both configuration and data. A GPIO pin provides flexible interrupt driven support through configurable mapping to event-triggers (i.e. reception of packet, TX FIFO underflow, etc). The TI CC11xL radio modules support a flexible packet structure, consisting of a preamble, a synchronisation word, a fixed/variable/infinite payload and a cyclic-redundancy check.

3.2.4.1 Wake-up Sequence Generation

The AS3930 wake-up sequence consists of three parts, namely a carrier burst, a preamble and a wake-up identifier. Figure 3.15 illustrates the structure of the wake-up sequence. The carrier burst consists of a sequence of alternating 1's and 0's at 125kHz for a duration of 2.4ms. The preamble consists of a 7-bit sequence (0101010) at a symbol rate of 2730 symbols/s. The wake-up identifier consists of the user-defined Manchester encoded 16-bit identifier (e.g. 0x96), also transmitted at a symbol rate of 2730 symbols/s.

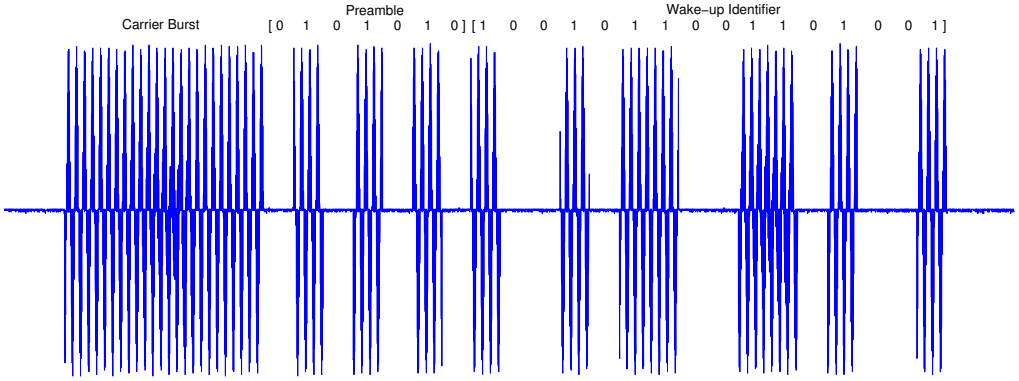

Figure 3.15: AS3930 wake-up sequence structure.

Figure 3.16 illustrates the signalling frequencies and their relations to the wake-up sequence. The TI CC110L radio module supports OOK modulation at a maximum bit rate of 250kbps. This means sending a single bit at 250kbps will induce a carrier frequency burst of 4us, which is conveniently half the period of the 125kHz square-wave signal needed for the wake-up sequence. Therefore, repeatedly sending '1' and '0' bits at 250kbps generates a waveform that has an envelope of a 125kHz square-wave.

The AS3930 supports a symbol rate of 2730 symbols/s to encode the preamble and wake-up identifier. The required 2730 symbols/s can be generated by sending a sequence of $\lceil \frac{250 \times 10^3}{2730} \rceil = 92$ symbols $= 11.5$ bytes at a bit rate of 250kbps. As the wake-up identifier is Manchester encoded, each 1-bit is represented as a binary 10-sequence and each 0-bit is represented as a binary 01-sequence, resulting in the transmission of 22 bytes at 250kbps per Manchester encoded bit. By configuring the preamble and sync words to 0xAA, utilising the infinite packet length feature, and disabling the CRC field in the CC110L packet structure, the wake-up sequence can be successfully generated through careful packet construction.

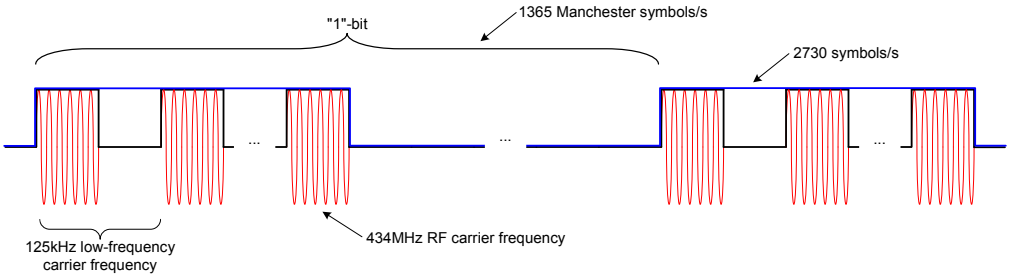

Figure 3.16: Signalling frequencies and their relation to the wake-up sequence.

The TI CC11xL radio modules supports a 64-byte transmission FIFO queue. When transmitting a packet longer than 64-bytes, as required for the wake-up sequence, it is imperative that the transmission FIFO queue is non-empty and does not overflow during transmission over the air [96]. An interrupt-driven approach ensures that the FIFO queue is sufficiently full, while a wake-up data structure maintains a state machine for determining the next byte to be written to the queue. The FIFO queue is initially filled with 64-bytes prior to RF transmission. During transmission, the CC110L issues an interrupt when the FIFO queue has reached a specified fill level (i.e. 53 bytes left in queue), and the FIFO is then re-filled. The FIFO is repeatedly re-filled until the complete wake-up sequence has been transmitted.

3.2.4.2 Measurement Data Transmission

Measurement data is communicated between the cluster node and the cluster-head at a data rate of 1.2kbps using OOK modulation. The measurement data rate is arbitrarily chosen for prototype implementation, and can be adjusted to suit the specific target application. The cluster-head microcontroller receives an interrupt from the TI CC110L radio module to indicate when a new packet arrives. Upon reception of the interrupt, the microcontroller checks the CRC field, extracts the fixed length payload and presents the measurement data to the data interface.

There is no per-packet flow control between the cluster-head and cluster node. This means there is no acknowledgement or retransmission primitives supported over the data channel. If measurement data packets are lost, it is assumed the data stream can be recovered through manual read out of the cluster node's non-volatile memory.

3.3 Sensing & Processing

3.3.1 Microcontroller

There is a vast array of commercial microcontrollers that satisfy the low power and peripheral requirements imposed by the cluster node and cluster-head. An extensive survey of microcontroller manufacturers resulted in a short list of three devices, as summarised in Appendix B. The final decision on which microcontroller to use considers multiple objectives, including functional aspects (sleep current, number of peripherals, GPIO port density, memory size, etc.) and non-functional aspects (development board availability, reference designs, IC package options, availability of tool chains, etc.). The Microchip PIC24FJ128GA310 family of 16-bit microcontrollers is chosen over the Atmel and the Texas Instruments alternatives, primarily due to its sleep performance and having prior development experience [64] on similar Microchip microcontrollers.

Although the PIC24FJ128GA310 microcontroller is well suited to the task at hand, it is over dimensioned with respect to its peripherals, and is restrictive on the availability of IC packages. This is further compounded by the unavailability of a commercial development board for this specific device. Therefore, the PIC24F32KA302 microcontroller, which is within the same PIC24F product family of 16-bit microcontrollers and has at least as good sleep performance, is chosen for prototype implementation. A commercial development board [87] for the 28-pin DIL package is leveraged for rapid prototype development.

The PIC24F32KA302 supports two SPI interfaces, two UART interfaces, an integrated real-time clock and calendar, three 16-bit timer modules, a 12-bit ADC, external interrupts and several GPIO pins. The microcontroller includes an internal 8MHz clock, and supports the addition of an external crystal for accurate real-time clock maintenance.

The Microchip PIC24F development tool chain consists of the MPLAB IDE v8.84 [88], integrated with an ANSI-compatible C compiler MPLAB C30 [86]. The C30 compiler optimises and compiles C source code into assembly. The generated assembly is then assembled with libraries, and linked together into a single Intel hex file for programming. The Microchip PICkit 3 [89] is used to program the PIC24F through a proprietary In-Circuit Serial Programming (ICSP) interface.

3.3.2 Power Management

The PIC24F32KA302 supports several sleep modes, with the most energy efficient being deep sleep mode [90]. The processor clock is turned off along with all peripherals during deep sleep, apart from the real-time clock (which is driven by an external 32.768kHz crystal) and an external interrupt pin. The contents of

all general purpose data registers and special function registers are lost during deep sleep, however, the state of all GPIO output latches are retained.

The sleep performance of the PIC24F32KA302 is of particular interest to the ultra-low power operation of the cluster node. Not only does the microcontroller's current drain go as low as 700nA @ 2.5V, but it can be woken-up by either a real-time clock alarm or an external interrupt trigger. These particular hardware features bode well with the reactive processing of the cluster node. In particular, the real-time clock alarm is aligned to the periodic measurement interval, and the external interrupt is driven by the wake-up interrupt signal from the AS3930 wake-up receiver. This ensures that the microcontroller can sleep for as long as possible, while still servicing all sources of internal and external stimuli.

Duty-cycling internal and external peripherals is an effective power management technique. This is achieved with the PIC24F32KA302 using two methods. The first is through the setting of peripheral module disable bits [90], which explicitly turn off all clock and power signals to unused internal PIC peripherals. The second method is through analog power switches digitally controlled through PIC GPIO pins. The TS5A3167 SPST analog switch from Texas Instruments [98] is used to power gate the data radio and the non-volatile memory. The quiescent current of the switch and any current leakage through the PIC GPIO pin is extremely low (i.e. on the order of 0.01μA).

3.3.3 Temperature Sensor

Thermistors (or thermally sensitive resistors) are passive devices which change their internal resistance according to the applied temperature [24]. There are two types of thermistors: positive temperature coefficient (PTC) and negative temperature coefficient (NTC) thermistors. The PTC and the NTC thermistors are differentiated by their respective positive or negative change in thermistor resistance as temperature increases. The NTC type thermistors have greater sensitivity than the PTC alternatives [24].

The 44006RC thermistor from Measurement Specialities [85] is selected for the cluster node design. The thermistor is an epoxy-coated NTC sensor which meets the stringent temperature accuracy requirements of at least ± 0.2°C over the majority of the operational temperature range. The non-linear transfer function of the 44006RC is given by the manufacturer in terms of a calibrated look-up table [85]. Figure 3.17(left) illustrates the calibration data together with a piecewise linear interpolation.

The temperature is measured by sampling the voltage across the thermistor when it is biased in a voltage divider circuit [8], as depicted in Figure 3.17(right). The PIC microcontroller samples the voltage across the thermistor using its internal 12-bit ADC, giving the value $X_{ADC} \in [0, 2^{12} - 1]$. Assuming no current

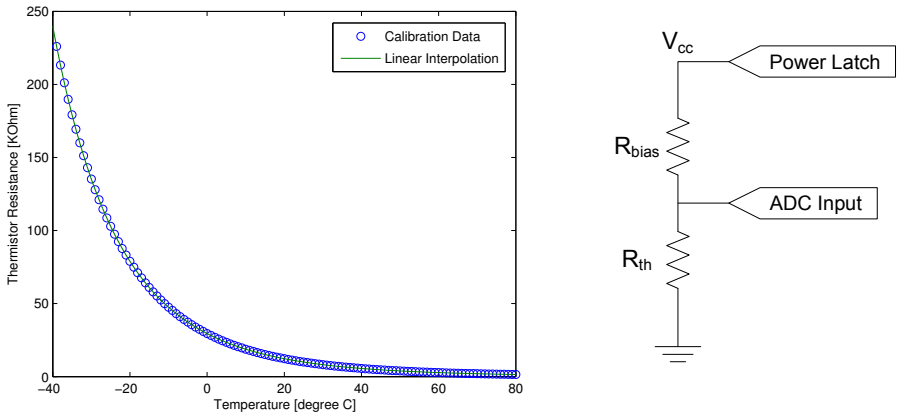

Figure 3.17: (left) Thermistor transfer function, and (right) thermistor measurement circuit.

leakage through the ADC, the voltage across the thermistor is given by V_{th} in Equation 3.2. Within a given sample resolution, one can calculate V_{th} from the digitised voltage measurement, and therefore express the thermistor resistance R_{th} in terms of X_{ADC} and the bias resistor R_{bias}, as in Equation 3.3. The measured temperature may then be calculated using a piece-wise linear interpolation based on the manufacturer-specified calibration data, as represented in Equation 3.4.

$$V_{th} = \left(\frac{R_{th}}{R_{th} + R_{bias}}\right) V_{cc} = \left(\frac{X_{ADC}}{2^{12}-1}\right) V_{cc} \tag{3.2}$$

$$\therefore R_{th} = \frac{X_{ADC} R_{bias}}{4095 - X_{ADC}} \tag{3.3}$$

$$Temperature = f_{linint}\left(R_{th}\right) = f_{linint}\left(\frac{X_{ADC} R_{bias}}{4095 - X_{ADC}}\right) \tag{3.4}$$

The bias resistor now needs to be selected such that the temperature resolution and accuracy requirements are satisfied. Selecting a large bias resistor would increase the dynamic range of the ADC, and thus improve the temperature resolution. However, thermistors are passive devices that dissipate heat when placed under load. The heat generated by using a large bias resistor would adversely effect the measured temperature, leading to self-heat error. The value of the bias resistor R_{bias} is analytically determined in terms of worst-case temperature resolution and worst-case self-heat error over the operational range of -40 to $+80°$C.

Due to the non-linearity of the thermistor's response, the biasing trade-off is analysed by calculating the proportion of the total operational temperature range which satisfies a worst-case temperature resolution ($\leqslant 0.2°C$) and a worst-case self-heat error ($\leqslant 0.05°C$). A Matlab simulation is used to evaluate the most suitable bias resistor for the thermistor measurement circuitry.

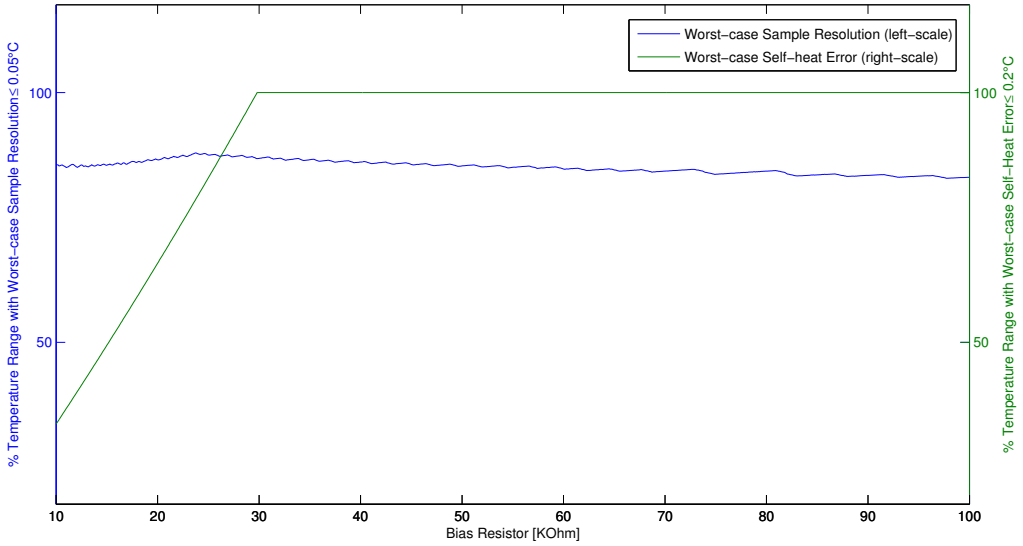

Figure 3.18: Evaluation of bias resistor based on worst-case temperature resolution and self-heat error.

The results in Figure 3.18 indicate that a bias resistor of approximately 29.9kΩ satisfies the self-heat error requirement over the entire operational temperature range. However, only 85.3% of the operational temperature range satisfies the temperature resolution requirement. Nevertheless, this biasing point is a good trade-off given the majority of the operational temperature range covered. Adjusting to standard component values, a bias resistor of $R_{bias} = 30\text{k}\Omega$ is chosen for the prototype implementation.

3.3.4 Software Architecture

The cluster node is a reactive system [47] with its behaviour completely defined in terms of four distinct software activities, namely configuration, measurement, communication and sleep. The configuration activity is responsible for gathering the cluster node system parameters prior to deployment, and is therefore only active once during the initial configuration of the cluster node. The measurement, communication and sleep activities however are executed throughout the operation lifetime of the cluster node. The measurement activity is responsi-

ble for performing periodic measurements and storing the result into the data store. When indicated from the wake-up receiver, the communication activity constructs measurement data from the data store and transmits the data to the cluster-head. The sleep activity is a low power state where no other functionality is implemented.

The four cluster node activities are triggered through internal and external events. These events are realised through hardware interrupts originating from PIC microcontroller's built-in or attached peripherals. As described in Section 3.3.2, the PIC24F32KA302 microcontroller features a deep sleep mode, which can be awoken by either a real-time clock alarm or an external interrupt. Table 3.1 details the trigger type, the mapped hardware interrupt source and whether deep sleep wake-up is supported for the configuration, measurement and communication activities.

Activity	**Trigger**	**Interrupt Source**	**Deep Sleep Awake**
Configuration	External	UART Data RX	No
Measurement	Internal	RTC Alarm	Yes
	Internal	ADC Finished	No
Communication	External	Wake-up Radio	Yes
	Internal	SPI Data Ready	No
	Internal	Timer Expired	No

Table 3.1: Summary of trigger type, interrupt source and deep sleep awake support for each cluster node activity.

As indicated in Table 3.1 the PIC microcontroller can be awoken from deep sleep independently by either the measurement or the communication activity. The RTC alarm provides a periodic trigger for the measurement activity, whereas the wake-up radio interrupt triggers the communication activity. The hardware support for awaking the microcontroller from deep sleep has a profound impact on the software architecture of the cluster node.

Modern sensor node operating systems (see [40] for a survey) typically maintain a software-based logical clock for scheduling activities in the future, with the aim of sleeping in between scheduled activities where possible. This software-based logical clock is driven by the sensor node's physical clock using one or more built-in timer peripherals. Maintaining the logical clock requires periodically waking from deep sleep, updating the logical clock, and re-entering deep sleep. This intermittent sleep behaviour has a noticeable impact on overall power dissipation of the microcontroller when sleep dominates activity schedule [36].

Since both measurement and communication activities can bring the cluster node's microcontroller out of deep sleep, there is no need to support a software-based logical system clock. The microcontroller's real-time clock becomes the

physical and logical system clock of the cluster node, thus removing any need for software-based time maintenance. This paves the way for a light-weight software architecture where excessive software abstractions are replaced with a focus on fast deep sleep invocation without unnecessary interruption.

The interrupt source summary in Table 3.1 infers that the three cluster node activities have hardware independence, as they each have a unique interrupt source. Furthermore, as the software functionality of each activity is unique with no chance of data deadlock, one can conclude the cluster node activities are also independent. In order to support the independent measurement and communication activities, the cluster node software architecture must provide concurrency. That is, the cluster node must be able to service multiple data flows simultaneously. For example, the cluster node may be taking a temperature measurement, while simultaneously servicing a wake-up request which triggers the transmission of a measurement packet over the data radio. Figure 3.19 illustrates the concurrency mechanism supported in the cluster node software architecture, which is an adaptation of the function-queue-scheduling architecture presented in [19], and the event-based programming model presented in [52].

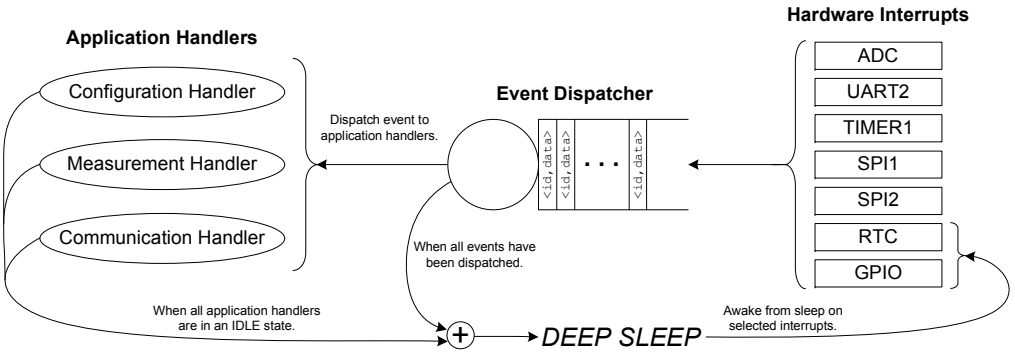

Figure 3.19: Execution flow through cluster node software architecture.

The key to the concurrency mechanism is the decoupling of interrupt-context and application execution using an event dispatcher. Every interrupt in the microcontroller is represented as an event tuple, consisting of a unique ID and some optional data associated with the interrupt. All events are inserted into a FIFO queue as they are triggered. The event dispatcher services the FIFO queue and routes each event to the appropriate application handler. The configuration, measurement and communication activities described earlier are implemented as application handlers in software (see Figure 3.20 for their respective StateCharts [26]). Concurrency is supported as time-critical hardware interrupts are serviced through interrupt service routines, while ensuring application handler code is executed to completion.

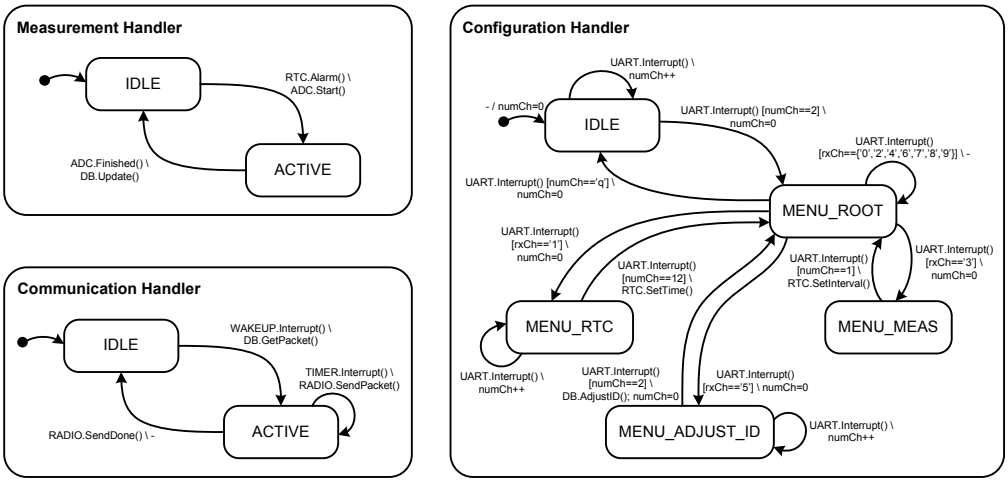

Figure 3.20: StateChart of measurement, communication and configuration handlers.

An advantage of this concurrency mechanism is the integration of deep sleep. In order to invoke deep sleep safely, that is all application handler variables are maintained throughout the low power mode, a single entry and exit point for deep sleep in the code must be ensured. If there are several deep sleep entry points, then the risk of corrupt state variables and the difficulty in debugging considerably increases. As illustrated in Figure 3.19, the invocation of deep sleep is dependent on two factors: all events must have been dispatched from the FIFO queue, and all application handlers must be in an idle state. Once these two conditions are satisfied, each application handler is requested sequentially to prepare for deep sleep by storing all necessary state variables to non-volatile memory (see Section 3.4.1 for details). Once completed, the microcontroller enters deep sleep. The microcontroller can then only awake from deep sleep through an RTC or external GPIO interrupt. When this occurs, each application handler is given the opportunity to restore their state variables from non-volatile memory before the event is dispatched from the FIFO queue.

The handling of built-in and attached peripherals is partitioned into low-level device drivers and modules. The device driver implements the interrupt service routine and primitive operations associated with the underlying hardware (e.g. SPI device driver). Modules are an encapsulation of one or many device drivers which combine to form a larger function (e.g. memory module). Figure 3.21 illustrates the cluster node software architecture consisting of application handlers, the event dispatcher, modules and device drivers.

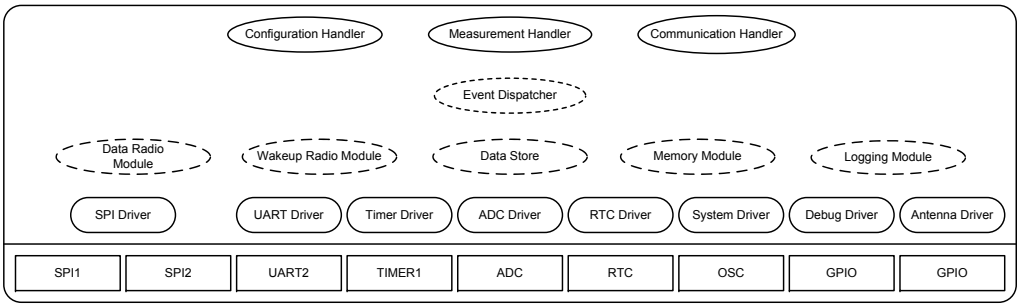

Figure 3.21: Software architecture of the cluster node.

3.4 Data Management

3.4.1 Non-volatile Memory

The cluster node requires non-volatile memory for the storage of measurement data and the preservation of local state variables during deep sleep. Several memory technologies are popular for integration with networked embedded systems, including EEPROM, Flash and Ferroelectric RAM (FRAM). Through an extensive component survey (see Appendix B for details), the FM25V10 FRAM memory device from Ramtron [93] is selected as the cluster node's non-volatile memory. The device has a superior energy efficiency compared with other memory technologies through a combination of low current drain and fast read/write access times.

The FM25V10 has 1Mbit of byte-addressable non-volatile memory. The PIC microcontroller interfaces to the non-volatile memory through a standard SPI interface. The FM25V10 also features a pre-configured 40-bit identification number unique to each device. This provides a unique identity for each cluster node without the need for extra hardware support or cumbersome software implementation.

3.4.2 Data Store

The purpose of the data store is to store and retrieve records of a specific structure. In the context of the cluster node, there are two types of records that must be stored and made available upon request, namely data records and time records. The record structure and how they are arranged in the non-volatile memory is detailed in the following section.

The data store is an overlay of the linear addressing space provided by the FM25V10 FRAM device. The physical address space of the FM25V10 FRAM

is byte-aligned, however to optimise the command structure on the SPI bus and to prevent unnecessary separation and concatenation of variables within records, a 16-bit word-aligned logical address space is constructed. Figure 3.22(a) illustrates the logical address space of the FRAM and the structure of the data store. The lowest 64 words are reserved for system parameters, specifically preserving important state variables during deep sleep. The remaining memory is used to store data and time records.

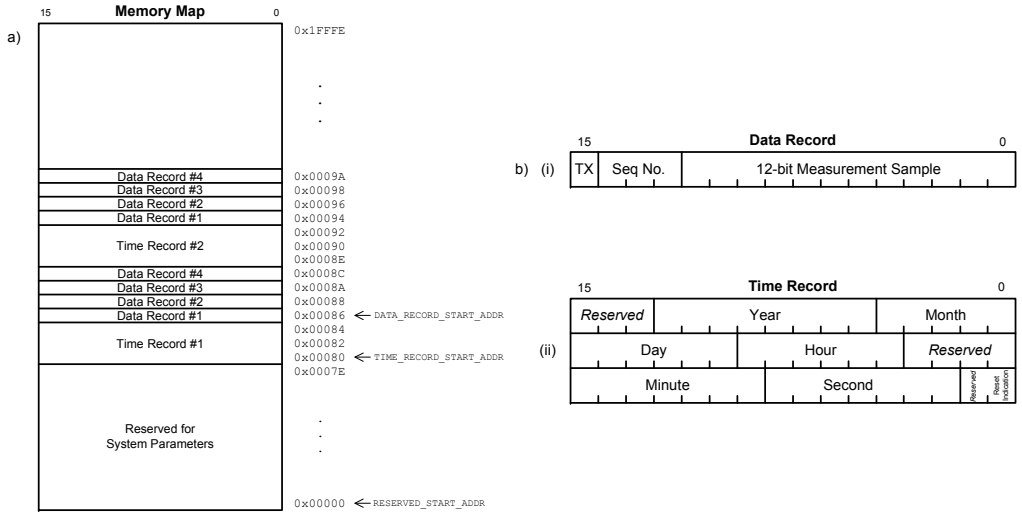

Figure 3.22: a) Memory map of the FM25V10 FRAM; b) (i) Data record structure, and (ii) time record structure.

The structure of the data and time records are illustrated in Figure 3.22(b). The data record contains the 12-bit measurement sample value from the ADC, a 3-bit sequence number and a bit field indicating if the record has already been transmitted. The time record contains the complete date and time in binary-coded-decimal format, with temporal resolution of one second. The values are extracted from special function registers maintained by the PIC real-time clock [90]. A reset indication bit field is used to indicate the first time record after a system condition. The remaining unused bit-fields are reserved for future use.

The target application requires the cluster node to sample the temperature sensor at periodic intervals. It is clear from this requirement that each temperature measurement must have a time associated with it, making it possible to reconstruct the time-series data from complete or only partial data availability at the cluster-head. However, it is not necessary to store a time record for every data record, but rather only store one time record per α data records. Since the time records are three times as large as the data records, the memory saving (i.e. increase in data lifetime) becomes significant with a suitable choice of α.

The ability to configure the number of data records stored per time record enables the data store to be dimensioned according to the supported hardware configuration and the deployment characteristics of the target application. The value of α is configured at compile-time, thereby fixing the data store structure for the operation lifetime of the cluster node.

Determining the most appropriate value for α depends on the intended data lifetime of the cluster node. That is, the number of years of periodic measurements the cluster node is expected to store in its data store. The cluster node data lifetime L_{data} can be evaluated using the analytical expression in Equation 3.5.

$$L_{data} = \frac{k_1 \left(M_{total} - M_{persistent}\right)}{\gamma \left(\lfloor \frac{S_{time}}{\alpha} \rfloor + S_{data} \right)} \tag{3.5}$$

Figure 3.23 illustrates the calculation of data lifetime L_{data} versus the number of measurements per hour γ, for several values of α. The calculations performed assume the total logical memory address space contains $M_{total} = 65536$ words with the first $M_{persistent} = 64$ words reserved for system variables, and the size of the data and time records are $S_{data} = 1$ and $S_{time} = 3$ words respectively. A constant scaling factor $k_1 = \frac{1}{24 \text{ hours} \times 365 \text{ days}}$ is used to ensure all input parameters are in the appropriate units. The results in Figure 3.23 indicate that a choice of $\alpha = 4$, that is one time record to every four data records, ensures the cluster node can comfortably satisfy the minimum operation requirement of 3 years with a 1 hour measurement interval.

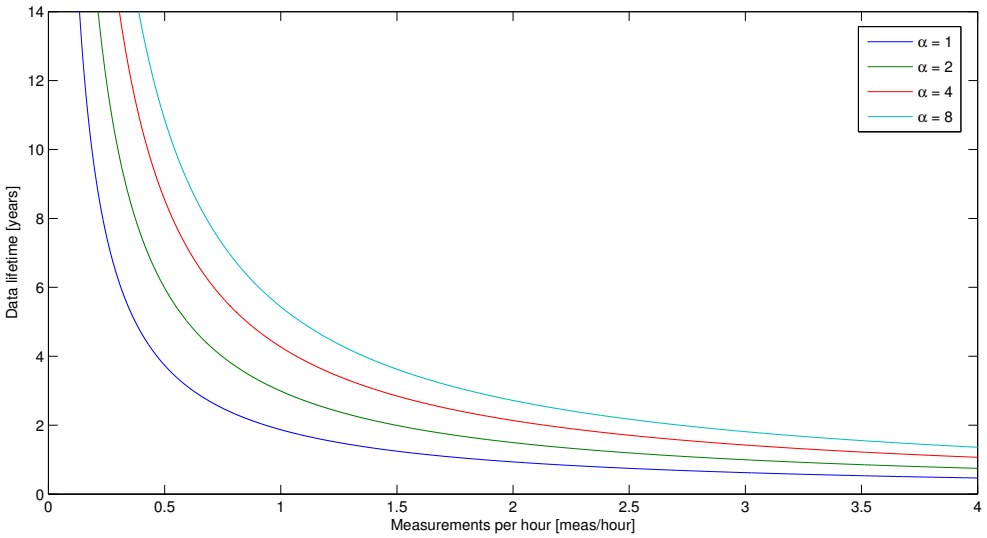

Figure 3.23: Data lifetime analysis.

3.4.3 Data Aggregation & Multi-user Access

The cluster-head aggregates measurement data from the cluster nodes using a pull-based data collection mechanism [15]. The data aggregation is facilitated by the wake-up radio coupled with a TDMA-based multi-user access scheme for orderly transmission of measurement data from each cluster node. Each node is allocated a finite time slot to transmit measurement data to the cluster-head, with the time slot being indexed by the cluster node's unique identity. As only 256 nodes are supported in a single cluster, the cluster node's 40-bit identity provided by the non-volatile memory (see Section 3.4.1) is truncated to 8-bits. The TDMA-slot time allocated per cluster node is only long enough to support a single measurement packet having a maximal length.

The pull-based data collection mechanism places responsibility on the cluster-head to request measurement data often enough, but not too often such that power is wasted through exorbitant wake-ups. For example, if a cluster node receives a wake-up request but does not have any measurement data to send, it saves energy by re-entering deep sleep, effectively wasting the allocated TDMA slot.

The structure of the measurement packet is depicted in Figure 3.24. The packet payload consists of the 8-bit cluster node identification, a time record, followed by up to N_{max} data records in chronological order. The time record is selected such that it is associated with the first data record in the packet payload. In the prototype implementation, $N_{max} = 10$ data records, which results in a TDMA-slot time of 260ms (assuming 1.2kbps data rate and the inclusion of a guard time).

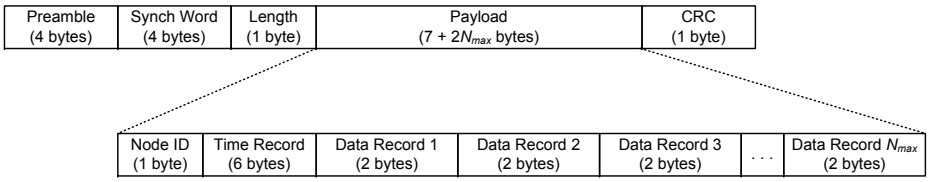

Figure 3.24: Structure of measurement packet.

The cluster-head does not implement any form of node-discovery. This simplifies node deployment and reduces implementation overhead at the cost of slightly increased idle listening at the cluster-head. However, since the number of cluster nodes (i.e. 256) and the worst-case TDMA-slot time (i.e. 260ms) are fixed, the worst-case cluster-head idle listening time is bound to approximately 67 seconds. Assuming a realistic deployment scenario (e.g. a measurement interval of 1 hour, requesting data every 3 hours), this translates to a cluster-head radio reception duty-cycle of $\frac{256 \times 0.260 \text{ seconds}}{10800 \text{ seconds}} \times 100\% \approx 0.62\%$. This is considered an acceptable overhead for a wireless data retrieval mechanism.

A single-pass algorithm is used to construct the measurement packet payload based on the time and data records stored in the cluster node's data store. The algorithm iterates through the data records in reverse-chronological order, and sequentially adds data records to the packet payload until a stopping condition is met. The algorithm has three stopping criteria; (i) if a total of N_{max} data records have been found, (ii) if a data record has already been transmitted as indicated by its TX bit field, or (iii) if an observed time record is the first after a system reset as indicated by its reset indication field. The justification for these stopping conditions are the following: (i) the packet payload only supports a maximum number of data records to bound TDMA-slot time, (ii) the cluster node only transmits measurement data that has not previously been transmitted to save precious energy reserves, and (iii) the time record's reset condition prevents the mixture of old and new data streams within a single measurement packet.

The algorithm used to construct a measurement packet payload from the cluster node's data store is exemplified in Figure 3.25. In the illustrated example, a total of eight valid data records are found until a previously transmitted record is identified. The time record associated with the earliest valid data record is included in the packet payload.

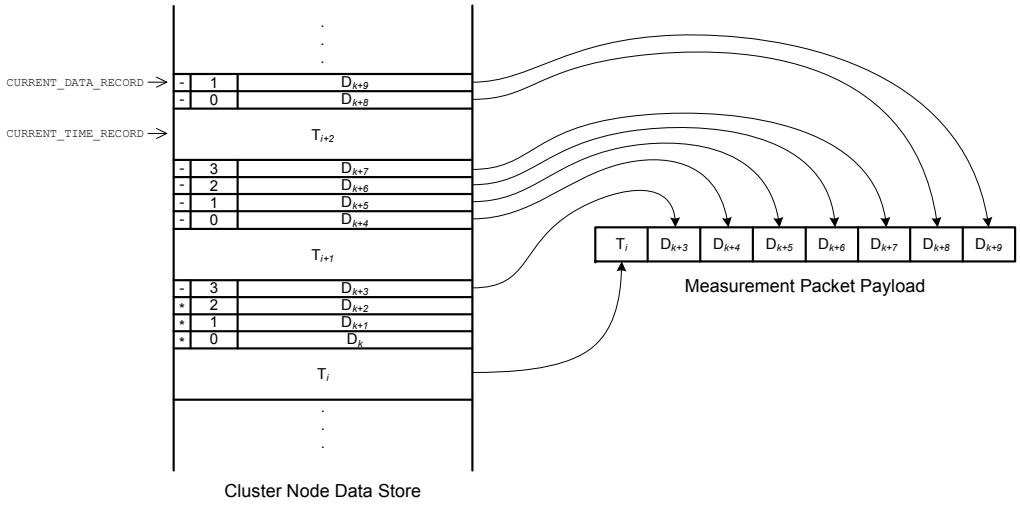

Figure 3.25: Example of packet payload construction from data store.

3.4.4 Configuration & Data Interfaces

The deployment of the wireless hierarchical sensor network is dependent on the availability of configuration and data interfaces. The configuration interface toward the cluster node is used for pre- and post-deployment tasks such as the setting of system time, initialisation of measurement interval and the download

of stored measurement data for data redundancy and validity purposes. The data interface toward the cluster-head is used for the presentation of measurement data received from the cluster in response to a measurement request. The configuration and data interfaces also serve as invaluable debug interfaces during prototype development and testing.

In an aim to reduce sensor network deployment barriers, the configuration interface should be easy to use, human-readable and should support cross-platform operation. An ASCII-based menu system is devised for the cluster node which operates using a serial over USB interface. The main menu of the prototype cluster node's configuration menu is listed in Figure 3.26(a). The ubiquity and portability of USB makes it well suited to the cluster node, while an ASCII-based menu system provides an easy to use and human-readable interface. Traversing the menu is performed in discrete ASCII-based commands, making it possible to automate node configuration with the use of scripts.

a)

```
+-----------------------------------------------+
| Wireless Hierarchical Sensor (Cluster node) |
+-----------------------------------------------+
|  0.  Download measurement data                |
|  1.  Set system clock                         |
|  2.  Display system clock                     |
|  3.  Set measurement interval                 |
|  4.  Display node information                 |
|  5.  Adjust node identification               |
|  6.  Display data radio settings              |
|  7.  Display wake-up radio settings           |
|  8.  Display memory settings                  |
|  9.  Erase all measurement data               |
+-----------------------------------------------+
 >
```

b)

```
+============================================================+
|           Wireless Hierarchical Sensor (Cluster-head)      |
+============================================================+
|      NODE        DATE          TIME         SEQ       VALUE     |
|============================================================|
|      0x4B     03/09/2012     16:16:10        0         0x232     |
|      0x4B     03/09/2012     16:16:10        1         0x233     |
|      0x4B     03/09/2012     16:16:10        2         0x233     |
|      0x4B     03/09/2012     16:16:10        3         0x232     |
|      0x4B     03/09/2012     16:16:10        0         0x232     |
|      0x4B     03/09/2012     16:16:10        1         0x232     |
|      0x4B     03/09/2012     16:16:10        2         0x232     |
|      0x4B     03/09/2012     16:16:10        3         0x231     |
|      0x4B     03/09/2012     16:16:10        0         0x233     |
|      0x4B     03/09/2012     16:16:10        1         0x232     |
+============================================================+
```

Figure 3.26: Listing of the prototype a) cluster node configuration and b) cluster-head data interfaces.

The purpose of the cluster-head data interface is more general, in that its implementation depends on the target application of the wireless hierarchical sensor network. If the cluster-head is a stand-alone unit, as is the prototype developed herein, a human-readable serial over USB interface would suffice (e.g. see Figure 3.26(b)). However, if the cluster-head is to be integrated with an existing wireless sensor node (e.g. a GPS-equipped wireless sensor node [34]), then perhaps a standard bus (e.g. SPI or I^2C) coupled a bi-directional command protocol would be more suitable.

3.5 Prototype Implementation

A prototype wireless hierarchical sensor network is illustrated in Figure 3.27. The prototype cluster-head and cluster node are based entirely on the detailed design presented in Sections 3.1 - 3.4. The cluster-head is implemented as a stand-alone unit, meaning the upper-tier multi-hop wireless network is not included in the prototype. The following sections annotate the prototype construction.

Figure 3.27: Final prototype of the ultra-low power wireless hierarchical sensor network, consisting of (left-to-right) the cluster-head, the OOK demodulator, and cluster node.

3.5.1 Hardware

The cluster node prototype is depicted in Figure 3.28. The PIC microcontroller, wake-up radio, data radio and auxiliary circuity are soldered onto a single piece of stripboard (also known as Veroboard). The stripboard is enclosed in a protective aluminium enclosure. A 6-pin DIN connector is used to interface a serial over USB cable [83] to the RS-232 line driver and SMA connectors are used to interface all RF interconnects.

The cluster-head prototype is illustrated in Figure 3.29. The data radio and the RS-232 line driver are both directly interfaced from the PIC development board [87]. The output of the RS-232 line driver connects to a serial over USB cable [83] via a 6-pin DIN connector. Panel mount switches and LEDs provide basic interactive features.

A design-for-test philosophy, based on the author's embedded system design experiences to date and those presented in [7], is applied to the construction of both cluster-head and cluster node hardware. The guiding principle is that prototype hardware platforms should be constructed in such a way that it is easy to test using a wide range of test instruments from its very inception. As the complexity of embedded systems increases, the dependency on a wide-range of

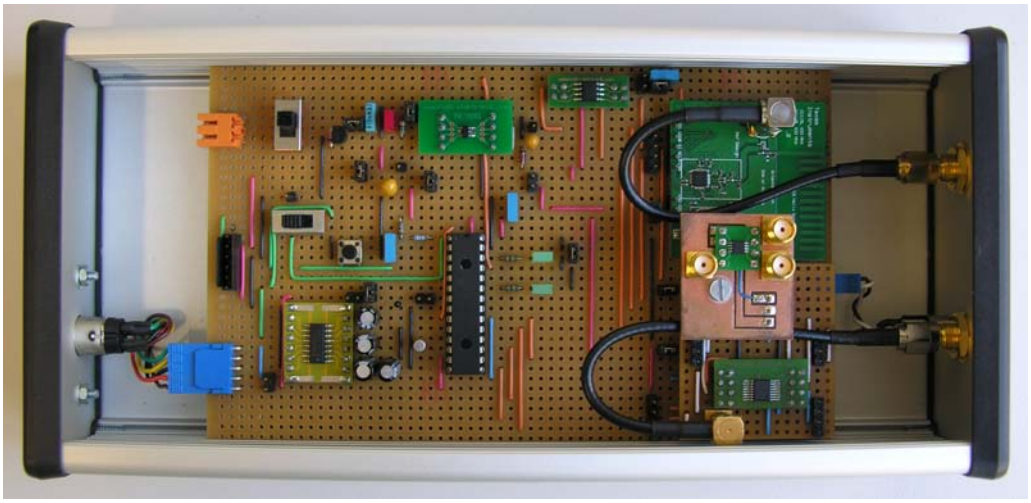

Figure 3.28: Cluster node prototype (top view).

debugging tools (e.g. multimeter, signal generator, oscilloscope, logic analyser, etc.) also increases. Incorporating this principle early in the hardware development cycle reduces overall development time; not only are complex hardware modifications late in the development phase avoided, but the ability to quickly set up complex test instruments to the prototype hardware promotes rapid software debugging. The following list summarises the most important techniques which were successfully applied during the construction of the prototypes:

- *Modularity*: Decomposing a large circuit into an interconnection of modular components reduces complexity in areas such as component layout, circuit analysis, and overall system diagnosis.
- *Stackable-boards*: Typically a small area of a larger prototype circuit will be subject to frequent change. Constructing these sections on small stackable circuit boards greatly simplifies circuit refinement.
- *Jumpering*: Providing isolation jumpers on all power junctions in the power supply and all IC positive supply rails simplifies basic diagnostic and power profiling procedures.
- *IC Sockets*: Utilising sockets or daughter-boards for all integrated components (through-hole and surface mount) ensure that any faulty or damaged components can be removed with minimal effort.
- *Debug Port*: The use of a generic debug port (consisting of LEDs, GPIOs, and a RS-232 serial port together with string and variable logging functions) is vital for embedded hardware and software development.

Figure 3.29: Cluster-head prototype (top view).

- *Power Pin-out*: A single jumper pin on the ground plane and on all voltage regulator outputs provides fast verification of the power plane using a multimeter.
- *Bus Pin-out*: Providing fundamental bus pins (i.e. for SPI and I^2C) through permanently soldered pins enables fast and reliable connectivity of test instruments.

The OOK demodulation circuit is depicted in Figure 3.30. Due to component and PCB parasitics, the RF circuit is constructed on a double-sided epoxy FR4 PCB using 0402 SMT component packaging (where possible). The circuit board of the OOK demodulator, and also the antenna switch circuit, are custom-made circuit boards, machined manually by the author using a Proxxon MF70 [71] miniature milling machine.

In an effort to reduce the impedance path to ground (simplify coupling to RF test instrumentation) and limiting external RF interferences, the OOK demodulator circuit is housed in a die-cast aluminium enclosure. The antenna input and the demodulated output are connected through SMA connectors, while a power jack supplies the DC bias voltage. It is the intention that the design of a multi-layer PCB will integrate the digital and analog circuitry of the cluster node onto a single board. For example, a 4-layer PCB would provide isolated digital ground, RF ground, positive power, and digital signal routing layers.

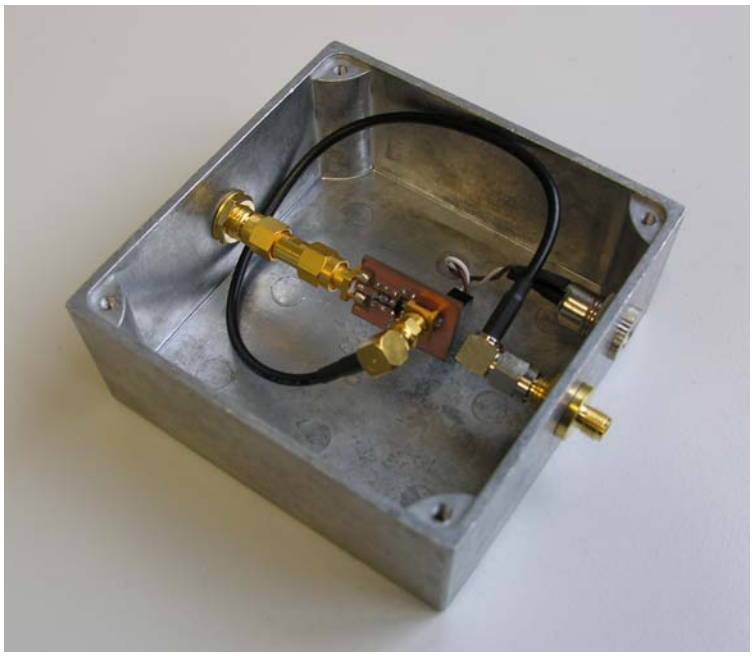

Figure 3.30: OOK demodulator prototype (top view).

3.5.2 Software

An adaptation of the rapid prototyping [21] software development process is applied to rapidly design, implement, build and test modular code segments aimed at a specific functional block (i.e. a software feature or a specific hardware block). Once a functional block is tested successfully in isolation, a snapshot of the complete code base is taken using a software versioning tool before continuing to build onto to the next functional block.

While iteratively combining functional blocks together into more complex functional blocks, there is also a need to work towards reducing uncertainty. That is to say, the selection of the next functional block must be such that it moves the development target closer to the next milestone. Aligning development milestones with the most difficult foreseeable technical challenge early in the prototyping phase will help to speed-up overall development. This strategy has been successfully demonstrated in the software development of the cluster node and cluster-head firmware. Figure 3.31 illustrates the software development cycle for the cluster-head and the cluster nodes.

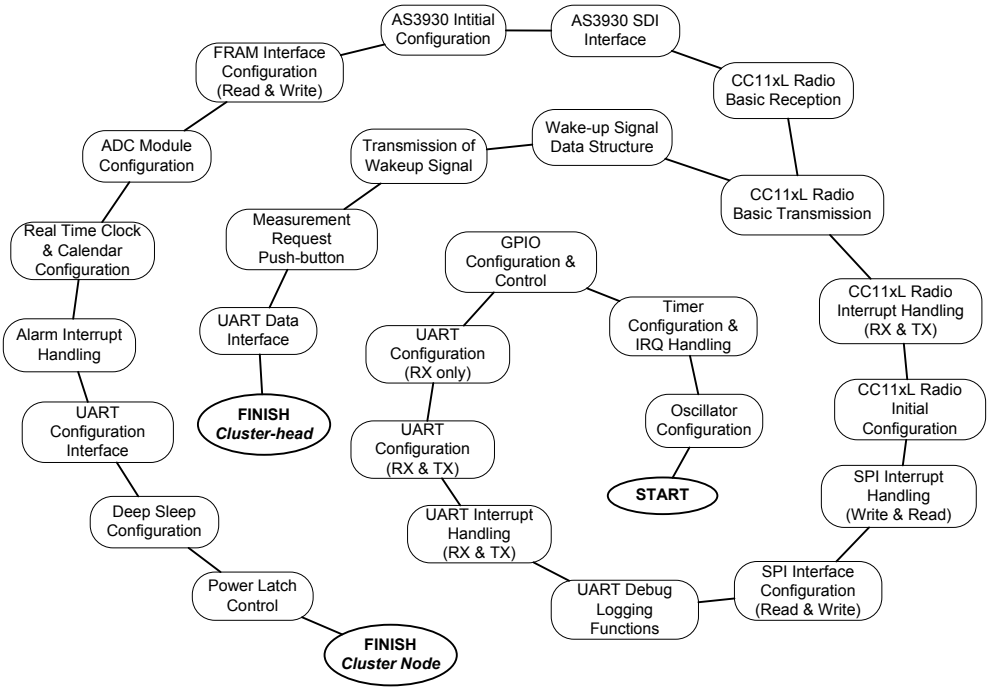

Figure 3.31: Software development cycle for the cluster-head and cluster node.

Chapter 4

System Evaluation

4.1 Power Performance

4.1.1 Power Dissipation

Given the resource and energy-constrained nature of wireless sensor networks, the power dissipation of a wireless sensor node is one of the most important performance metrics. Assuming the wireless sensor node is powered by a battery (e.g. [65]), or through energy scavenging techniques (e.g. [66]), the total power dissipation directly impacts the operation lifetime of the device in the field.

It is the intention of the wireless hierarchical sensor network to maintain several years of operation in the field. Specifically, the cluster node is intended to be deployed for several years powered by a single coin cell battery, whereas the cluster-head is intended to be powered by a higher capacity battery with the potential for recharge using energy harvesting techniques. The following section will evaluate the power dissipation of the cluster node and then use these measurements to estimate the operation lifetime of the cluster node and cluster-head.

The sources of power dissipation may be classified into two main categories, namely dynamic and static power dissipation [16]. The dynamic power dissipation of a device is dependent on the active tasks performed by the device (e.g. execution of the software state machine, activation of peripherals, response to external stimuli, etc.). The static contribution is the power dissipation independent of the device's active tasks (e.g. current leakage, bias current, regulator quiescent current, etc.). In order to investigate static and dynamic power dissipation contributions, the total power dissipation of the device under test is measured over a finite time interval, which is termed a power profile.

The power profile of the cluster node is measured using an Agilent N6705A DC Power Analyzer, with the experimental set up illustrated in Figure 4.1. The cluster node is the device under test, with the Agilent N6705A configured as a battery emulator supplying 3.0V to the voltage regulator of the cluster node. The

cluster node software is customised through build-flags to enable automatic node configuration and the removal of all debug features. The automatic measurement configuration erases all records in the data store, configures the real-time clock with a valid time and sets a 10 second periodic measurement interval. All debug LEDs, UART debug logs and the serial over USB interface are disabled so not to impact the power profile.

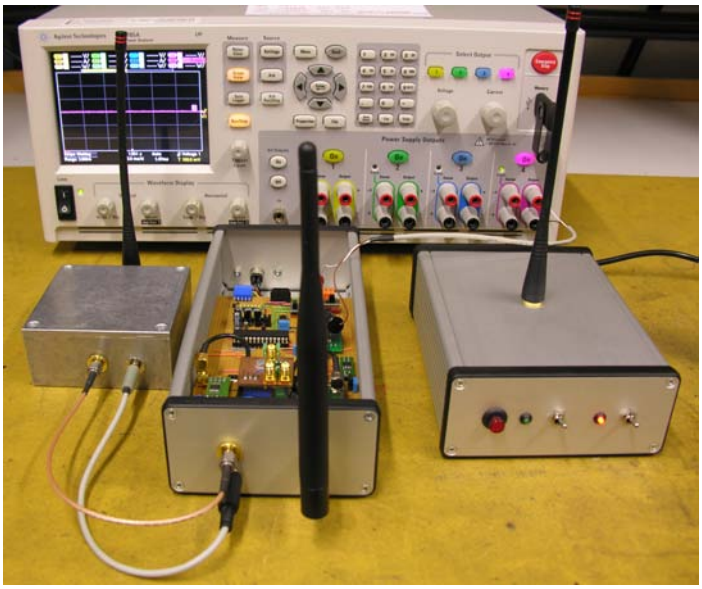

Figure 4.1: Experimental set up of cluster node to the Agilent N6705A DC Power Analyzer for power dissipation analysis.

Figure 4.2 illustrates the cluster node power profile with the complete prototype hardware, executing a test case of 10 temperature measurements followed by a radio wake-up and the transmission of the 10 previously stored measurements. The test case sequence is as follows: (i) the cluster node is powered on and auto-configured, (ii) the node enters deep sleep until it is woken by the 10 second periodic measurement alarm, (iii) the temperature sensor is powered on, sampled by the ADC before being turned off, the measurement is stored in the data store and (iv) the cluster node re-enters deep sleep. A total of 10 measurements at a 10 second interval are performed before (v) a wake-up sequence is received from the cluster-head and a measurement packet containing the past 10 data records is extracted from the data store and transmitted to the cluster-head.

The power profile in Figure 4.2 illustrates the relative power dissipation magnitude and time-scale of the configuration, measurement, communication, and sleep activities. The configuration activity is only expected to be performed once prior to deployment. Since the cluster node can be powered through the serial to USB configuration interface, the power consumed by this activity does not im-

pact the node's operation lifetime. Irrespective of the cluster node's deployment configuration, the combined execution time of measurement and communication activities are several orders of magnitude shorter than the time spent sleeping. This suggests the extended operation lifetime can be achieved through optimisation of the static power dissipation contributors.

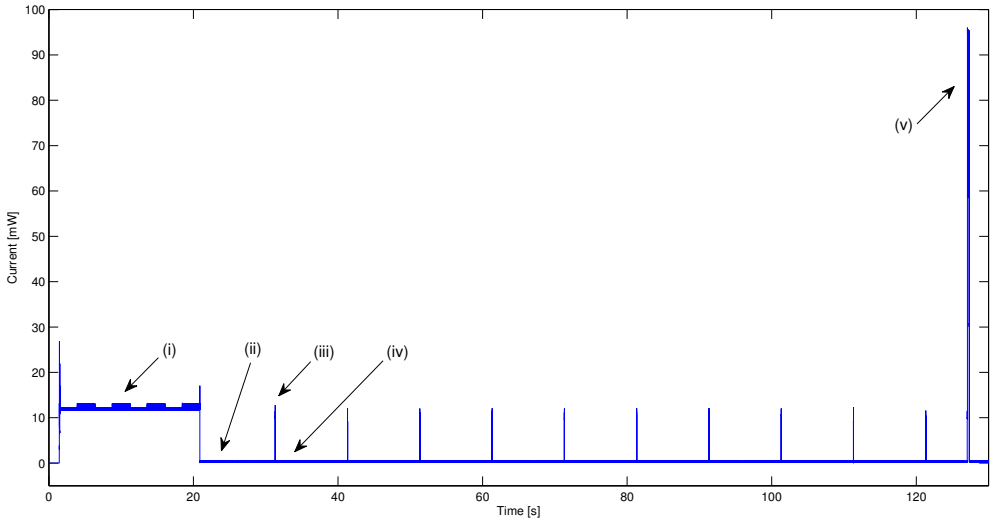

Figure 4.2: Cluster node power profile including antenna switch and OOK demodulator DC biasing, measured at 3.0V.

The measurement of static and dynamic power dissipation contributors in isolation is difficult. This is mainly attributed to the inability to isolate static and dynamic power dissipation contributions from within a single component. For example, the quiescent current of the cluster node's low drop out (LDO) voltage regulator [95] is dependent on the instantaneous load, which is in turn dependent on both static and dynamic components. However, one may make an approximation of the static power dissipation with knowledge of the circuit schematics and component data sheets. Table 4.1 summarises the cluster node's static power dissipation contributors.

It is clear from Table 4.1 that the antenna switch (see Section 3.2.1) and the OOK demodulator DC biasing (see Section 3.2.2.3) are the largest two contributors, by at least a factor of two. However, according to its datasheet [98], the analog switch has a quiescent current drain of 0.01μA. This large discrepancy is attributed to the impedance mismatching of the device's input RF ports. The absorptive switch draws a higher current as the microstrip transmission lines [23] on the prototype PCB are not ideal. The OOK demodulator circuit on the other hand requires external biasing to ensure operation with small input signals. However, as detailed in Section 3.2.2.5, a power-optimised OOK demodulator using zero bias Schottky diodes [78] eliminates the need for this DC biasing.

Component	Static Power Dissipation [mW]
TS5A3167 Power Latch	0.12
AS3930 Wake-up Receiver	8.25
TPS76925 2.5V LDO Regulator	51
ADG918 Antenna Switch	120
OOK Demodulator DC Bias	129

Table 4.1: Summary of cluster node static power dissipation contributors.

Figure 4.3 illustrates a second power profile of the cluster node, with the antenna switch removed (emulating a final PCB with matching input RF ports) and using the power-optimised OOK demodulation circuit (i.e. eliminating external DC biasing). The test case sequence is the same as in the power profile of Figure 4.2, with the exception of only performing 3 measurements instead of 10. The power is expressed in a logarithmic scale to highlight the small power dissipation contributions during deep sleep.

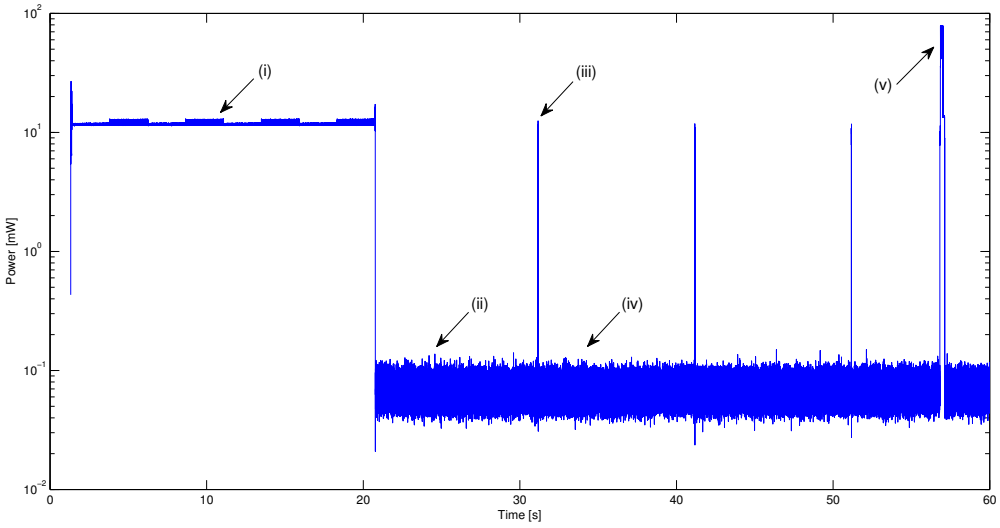

Figure 4.3: Cluster node power profile excluding the antenna switch and with the power-optimised OOK demodulator, measured at 3.0V.

The power profile in Figure 4.3 highlights the very low deep sleep power dissipation of the cluster node's PIC microcontroller. The average power dissipation during deep sleep is approximately 75μW (i.e. 25μA @ 3.0V). Removing the quiescent current contributed by the low drop out voltage regulator, the PIC microcontroller is only dissipating on average approximately 28μW (i.e. 9.3μA

@ 3.0V). This is exceptionally low, considering it is maintaining a real-time clock and is able to awake from deep sleep using a pre-configured alarm (i.e. the measurement alarm) and an external interrupt trigger (i.e. the wake-up interrupt).

The cluster node prototype board features power isolation jumpers (see Section 3.5.1 for details) which enable current drain monitoring of several components in isolation. This made it possible to perform a detailed component-level power dissipation analysis. Figure 4.4 illustrates the individual power dissipation of the microcontroller, data radio, non-volatile memory and wake-up radio. The test sequence consists of three measurements at a 10 second interval, followed by the reception of a wake-up sequence and transmission of a data measurement.

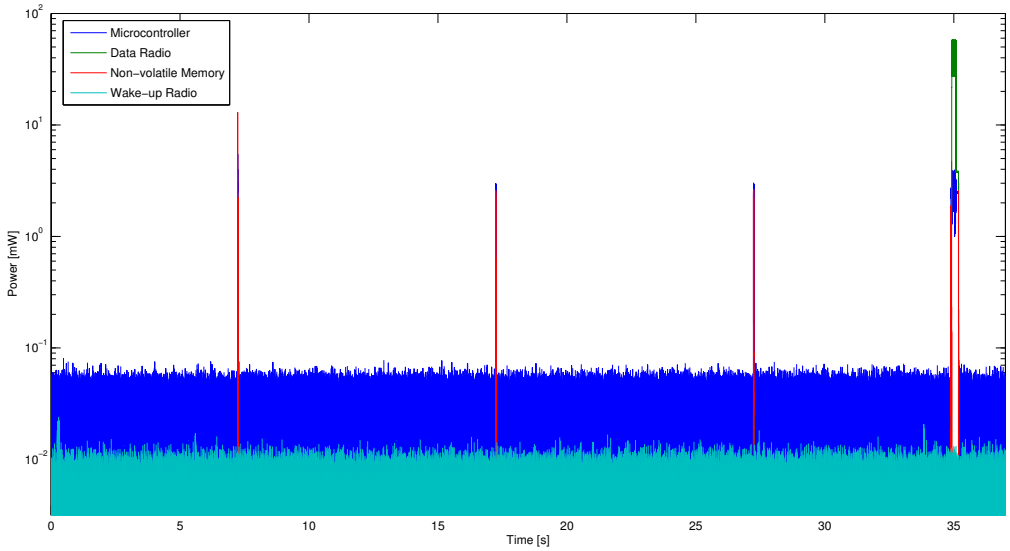

Figure 4.4: Component-level power dissipation of the cluster node's microcontroller, data radio, non-volatile memory and wake-up radio, measured at 2.5V.

The component-level power dissipation analysis in Figure 4.4 highlights the ultra-low power operation of the cluster node. All components are only powered when they are needed for measurement and/or communication activities, while the power dissipation of the wake-up radio and microcontroller components during sleep is stable and extremely low.

It is difficult to benchmark the power performance of the cluster node due to the lack of comparable wireless sensor nodes in the literature to date. However, to verify the ultra-low power criteria, it is of interest to evaluate the total average power dissipation at the cluster node's minimum functional requirement (see Section 2.5). Using the precision measurements captured in Figure 4.4, the duration and average power dissipation of each component is evaluated based on the cluster node's three primitive activities, as summarised in Table 4.2. The

use case considered assumes a temperature measurement once per hour, with a data request once every 3 hours. The analysis in Table 4.2 concludes that the cluster node has a remarkable 78.5μW total average power dissipation under the considered use case.

Activity	Component	Duration [s]	Power [mW]
Measuring	Microcontroller	0.0456	2.4399
	Memory		0.1811
	Wake-up Radio		0.0032
Communicating	Microcontroller	0.3020	2.5092
	Data Radio		22.1637
	Memory		0.0195
	Wake-up Radio		0.0032
Sleeping	Microcontroller	10799.6524	0.0236
	Wake-up Radio		0.0032
	LDO Regulator	10800	0.0510
		Total	**0.0785mW**

Table 4.2: Total average power dissipation of cluster node assuming a 1 hour measurement interval, with a data request every 3 hours. All components measured at 2.5V, except for the LDO regulator at 3.0V. The duration and average power dissipation values are extracted from Figure 4.4.

4.1.2 Operation Lifetime

The operation lifetime of the cluster node and cluster-head depends on both application-specific and network-specific configuration parameters of the wireless hierarchical sensor network. These parameters may be statically defined (i.e. at compile-time or during pre-deployment configuration) for the lifetime of the network deployment, or may be changed dynamically during deployment based on application-specific requirements. Table 4.3 summarises the application and network-specific parameters for a wireless hierarchical sensor network deployment.

The operation lifetime of the cluster node can be estimated analytically by considering the application and network-specific configuration parameters together with the total current drain of the cluster nodes three principle activities; measuring, communicating and sleeping. The configuration activity is not considered in the operation lifetime estimation as it is assumed that the cluster node is powered by an external power supply during configuration (e.g. through the serial over USB cable). Using the power profiles detailed in Section 4.1.1, the current drain and duration of the cluster node's principle activities can be conservatively approximated using piecewise linear functions. Figure 4.5 illustrates an

Parameter		Type	Configured
Data requests per hour	θ_r	Application	Deployment
Measurements per hour	θ_m	Application	Pre-deployment
Maximum data records per packet	N_{max}	Network	Compile-time
Maximum nodes per cluster	N_n	Network	Compile-time
Transmission slot duration	Δ_t	Network	Compile-time

Table 4.3: Summary of application and network-specific configuration parameters of a wireless hierarchical sensor network.

ADC measurement and the data radio sequence extracted from the power profile in Figure 4.3 expressed in milliamps. These two current waveforms are used to build the piecewise linear approximations. The ADC measurement waveform is independent of the measurement interval. However, the data radio waveform is dependent on the number of data records transmitted and the data rate in which they are transmitted. The current drain during sleep (including microcontroller and low drop out regulator) is considered constant, and is approximated by an average of 25μA @ 3.0V, as evaluated in Section 4.1.1.

The cluster node's operation lifetime L^{node} is estimated using Equation 4.1. The operation lifetime is estimated by dividing the available battery capacity C_{node} in mAh scaled by a discharge efficiency β, divided by the hourly current drain contributions for measurement, communication and sleep activities. The constant scaling factor k_2 ensures all input parameters are in the appropriate units. A description of all the input parameters is listed in Table 4.4.

$$L^{node} = \frac{k_2 \beta C_{node}}{\theta_m T_m I_m + \theta_r \sum_{i \in a,b,c,d} T_i I_i + \left(3600 - \theta_m T_m - \theta_r \sum_{i \in a,b,c,d} T_i\right) I_s} \tag{4.1}$$

In a real-world deployment of a wireless hierarchical sensor network, the cluster-head must be able to request data when there is data available in the cluster. This implies that the operation lifetime of the cluster node defines the minimum lifetime of the cluster-head. It is therefore interesting to quantify the operation lifetime of the cluster-head to ensure the wireless hierarchical sensor network can achieve a desired operation lifetime. However, it is difficult to accurately quantify the cluster-head's operation lifetime due to the vast array of hardware and software configurations that can provide the upper-tier multi-hop network connectivity. These additional hardware and software components will have a large impact on the overall power dissipation, and hence the operation lifetime of the cluster-head.

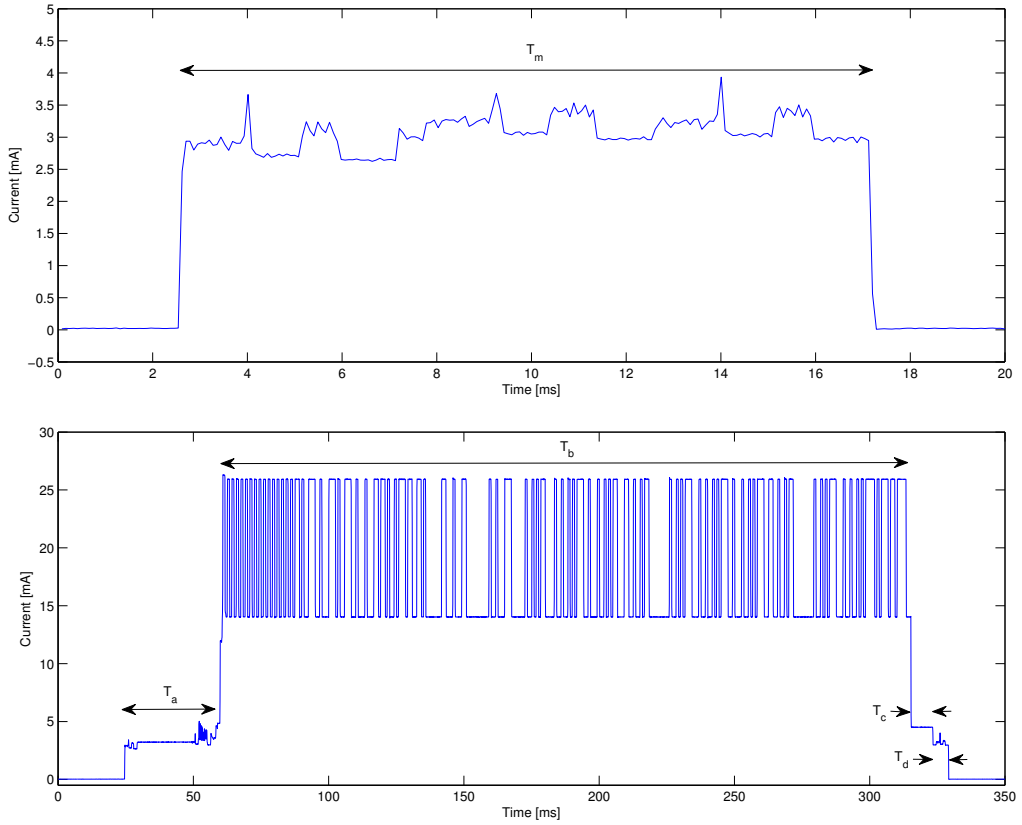

Figure 4.5: Measurement and communication waveforms used to construct piecewise linear approximation for operation lifetime analysis.

However, an optimistic estimation of the cluster-head operation lifetime is achievable by only considering the cluster-head's total radio activity, and assume it is in a low power sleep mode otherwise. The cluster-head must perform two radio activities: the transmission of the wake-up sequence and the reception of data measurements over a finite set of TDMA-slots. These two activities are triggered by the data request rate, as determined by the target application requirements. Using a similar approach as the one used with the cluster node, the operation lifetime of the cluster-head L^{head} is estimated using Equation 4.2.

$$L^{head} = \frac{k_2 \beta C_{head}}{\theta_r \left(T_w I_{TX} + \Delta_t N_n I_{RX}\right) + \left(3600 - \theta_r \left(T_w - \Delta_t N_n\right)\right) I_s} \tag{4.2}$$

Table 4.4 summarises the descriptions for all input parameters used in Equations 4.1 and 4.2, and specifies the values corresponding to the functional prototypes detailed in Section 3.5. Figures 4.6 and 4.7 illustrate the operation lifetime estimate based on the prototype cluster node and cluster-head respectively.

Parameter		**Value**
Measurements per hour	θ_m	360, 60, 6, 1, $\frac{1}{2}$, $\frac{1}{3}$ per hour
Data requests per hour	θ_r	$\frac{1}{3}$, $\frac{1}{2}$, 1, 2, 3 per hour
Battery capacity cluster node	C_{node}	1000 mAh
Battery capacity cluster-head	C_{head}	14000 mAh
Battery discharge efficiency	β	75%
Conversion constant	k_2	$\frac{60 \times 60 \text{ seconds/hour}}{24 \text{ hours} \times 365 \text{ days}}$
Maximum data records per packet	N_{max}	10 data records
Number of data records per packet	N_{data}	N_{max} if $\lfloor \frac{\theta_m}{\theta_r} \rfloor \geq N_{max}$ 1 if $\lfloor \frac{\theta_m}{\theta_r} \rfloor = 0$ $\lfloor \frac{\theta_m}{\theta_r} \rfloor$ otherwise
Maximum nodes per cluster	N_n	256 nodes
Transmission data rate	D	1200 bps
Transmission slot duration	Δ_t	260×10^{-3} seconds
Measurement activity duration	T_m	15×10^{-3} seconds
Measurement activity current drain	I_m	3.4 mA
Radio activity duration	T_a	34×10^{-3} seconds
Radio activity current drain	I_a	3.3 mA
Radio activity duration	T_b	$\frac{136+16N_{data}}{D} + 10\times10^{-3}$ seconds
Radio activity current drain	I_b	26 mA
Radio activity duration	T_c	$270\times10^{-3} - T_b$ seconds
Radio activity current drain	I_c	4.6 mA
Radio activity duration	T_d	6×10^{-3} seconds
Radio activity current drain	I_d	3.3 mA
Sleep activity current drain	I_s	25×10^{-3} mA
Wake-up sequence duration	T_w	17×10^{-3} seconds
Radio reception current drain	I_{RX}	16 mA
Radio transmission current drain	I_{TX}	29.2 mA

Table 4.4: Summary of input parameters and values used to evaluate the cluster node and cluster-head operation lifetime estimates.

The cluster node is able to achieve the functional requirement of 3 years at a measurement interval of 3 hours, as illustrated in Figure 4.6. In fact, the evaluation indicates more than 3 years operation lifetime is achievable with a data request rate as low as 2 requests per hour. There exist many other operation

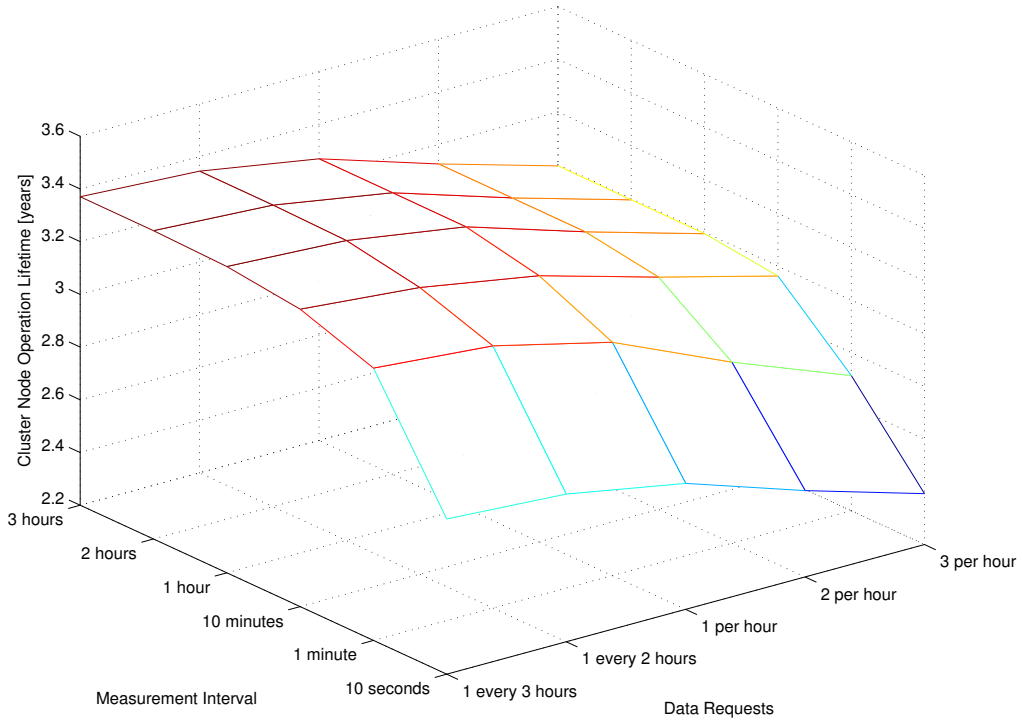

Figure 4.6: Cluster node operation lifetime estimate based on parameters listed in Table 4.4.

points (i.e. combination of measurement interval and data request rate) where the prototype cluster node supports greater than 3 years operation lifetime, however it may not be possible to retrieve all the data stored on the cluster node (e.g. 10 minute data rate, data request every 3 hours, with $N_{data} = 10$ records and $D = 1.2$kbps).

The operation lifetime significantly reduces in Figure 4.6 as the measurement interval increases from 10 minutes to 10 seconds. This is due to the measurement activity becoming the dominant factor in the power budget. Increasing the data request rate from 1 per hour to 3 per hour also reduces the operation lifetime, but to a lower degree, since many data records can be combined into a single data measurement packet.

The operation lifetime of the cluster-head is only dependent on the data request rate, assuming all other network-specific configuration parameters (see Table 4.3) are fixed prior to deployment. As illustrated in Figure 4.7, the evaluation indicates that the operation lifetime is significantly lower than the target of 3 years at data request rates higher than 1 per hour. This highlights the need for careful power analysis of both the cluster node and cluster-head in order to find

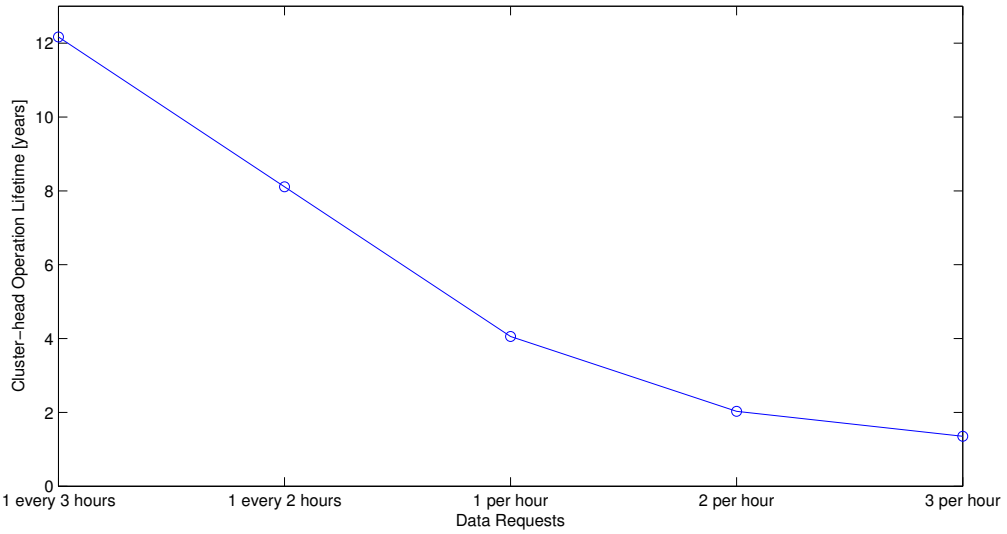

Figure 4.7: Cluster-head operation lifetime estimate based on parameters listed in Table 4.4.

the operation point where the wireless hierarchical sensor network can support a target application deployment lifetime.

As indicated in Figure 4.6, there are operation points defined by the application-specific configuration parameters where it is not possible to retrieve all data from the cluster. However, it is possible to increase the throughput of the wireless hierarchical network with negligible impact on the operation lifetime of cluster node and cluster-head. The evaluation presented in Figure 4.6 is based on a 1.2kbps data channel as implemented in the functional prototypes. However, according to the datasheet of the data radio [96], increasing the data rate up to 250kbps has a negligible impact on the transmission (I_{TX}) and reception (I_{RX}) current drain. It is assumed that increasing the data rate will not adversely affect the performance of the data channel due to the low range considered, coupled with the high data radio receiver sensitivity.

The implication of this is that the cluster node can send more data records per data measurement packet (i.e. increase network throughput for all nodes in the cluster) with negligible impact on the operation lifetime of the cluster node. Furthermore, since the TDMA-slot duration remains unchanged, there is negligible impact to the operation lifetime of the cluster-head. For example, increasing the data rate to $D = 250$kbps, and increasing the maximum number of data records per packet $N_{max} = 360$, more than 3 years of operation lifetime is supported on the cluster node at a measurement interval of 10 minutes, with the successful retrieval of all measurement data.

4.2 Communication Performance

4.2.1 Antenna & Antenna Switch

The performance of RF components and networks are typically quantified in terms of their efficiency to transfer power between input and output RF ports. A common method to quantify the performance is to measure the power reflected back from an input port when matched with a characteristic load impedance. The resulting metric is termed the reflection coefficient and is typically expressed in dB as the return loss. The return loss is also known as the S_{11} parameter, which is part of a more general framework of S-parameters (or scattering parameters) [11].

A variety of 434MHz compatible antennas are evaluated for the prototype wireless hierarchical sensor network. The return loss of the tested antennas varies significantly depending on the type of antenna, the antenna manufacturer and the type of ground plane coupled to the antenna. Figure 4.8(left) illustrates the return loss of a monopole, dipole and helical antenna, both with and without a ground plane. As the analysis indicates, the return loss varies significantly in both bandwidth and magnitude. The monopole antenna with a ground plane features a low return loss at the 434MHz carrier frequency while providing a sufficiently large bandwidth, and it therefore selected for the prototype implementation.

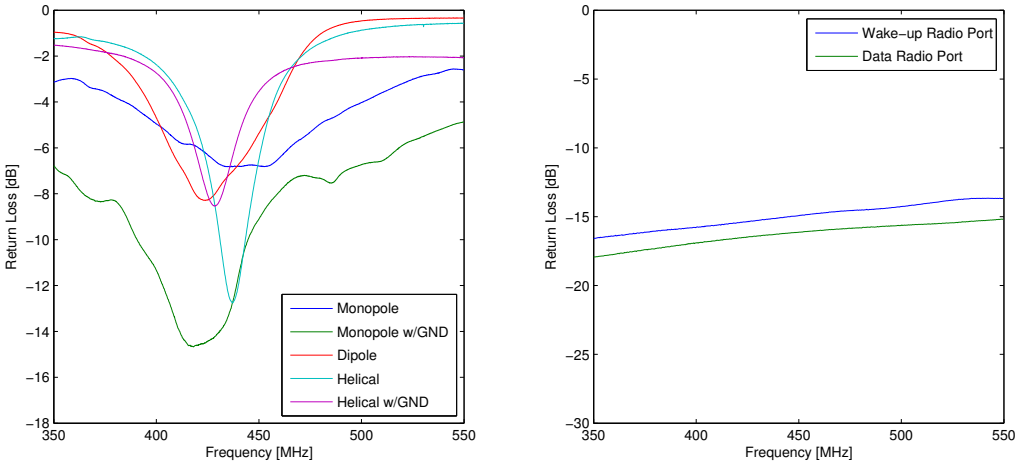

Figure 4.8: (left) Return loss of monopole, dipole and helical 434MHz antenna, and (right) return loss of antenna switch.

The return loss analysis in Figure 4.8(left) stresses the importance of an antenna ground plane, and highlights the grounding limitations in the prototype cluster node implementation. For example, the stripboard used to implement

the prototype cluster node circuit (see Section 3.5 for details) is well suited to prototyping digital or low frequency analog circuits. However, constructing high-frequency circuits or implementing analog and mixed signal circuits on stripboard is prohibitive. The cluster node stripboard contains several digital grounds which form high impedance paths to ground, severely affecting the performance of RF circuits such as the OOK demodulator. A solution enabling rapid-prototyping is to place the RF components into a diecast aluminium enclosure (see Section 3.5.1). An alternative is to design a multi-layer PCB incorporating suitable digital and analog ground layers.

Figure 4.8(right) illustrates the return loss measurements of the antenna switch. The measurements are performed by connecting 50Ω SMA terminators to the wake-up and data radio ports. The network analyser is then connected to the antenna port and measures the return loss through the antenna switch. Experiments indicate the 50Ω terminators have a return loss in the order of −40dBm, however the return loss of the terminator through the antenna switch is found to be less than half this optimal value.

The impedance mismatching on the antenna switch is attributed to the sub-optimal PCB. The microstrip transmission lines for each RF port on the PCB are not precisely matched at 50Ω, primarily due to the daughter board used for the ADG918 antenna switch integrated circuit. It is expected that a multi-layer PCB would significantly improve the RF performance of the antenna switch.

4.2.2 Wake-up Radio

4.2.2.1 OOK Demodulator

The RF performance of the OOK demodulator was evaluated in a similar way to the antenna and antenna switch in Section 4.2.1, that is measuring the return loss over the frequency range of interest. The return loss of the OOK demodulator circuit was measured using an Agilent 8720D Network Analyzer as depicted in Figure 4.9. The network analyser was configured to sweep over the 350MHz to 550MHz frequency range and over a range of input powers.

It is the intention for the OOK demodulator circuit to absorb as much incident signal power as possible at the 434MHz carrier frequency. As illustrated in Figure 4.10, the profile of the return loss varies significantly with input power resulting in a return loss between −5.8dB and −16dB. The dependence on input signal power is due to the non-linearity of the Schottky diodes in the voltage doubler circuit (see Section 3.2.2.3 for details). The parasitic resistance of the Schottky diodes changes with the incident signal power, which induces a change in the load resistance and reactance of the impedance matching circuit (see Section 3.2.2.4). This results in sub-optimal matching of the source and load impedances, and thus adversely affects the OOK demodulator return loss. The

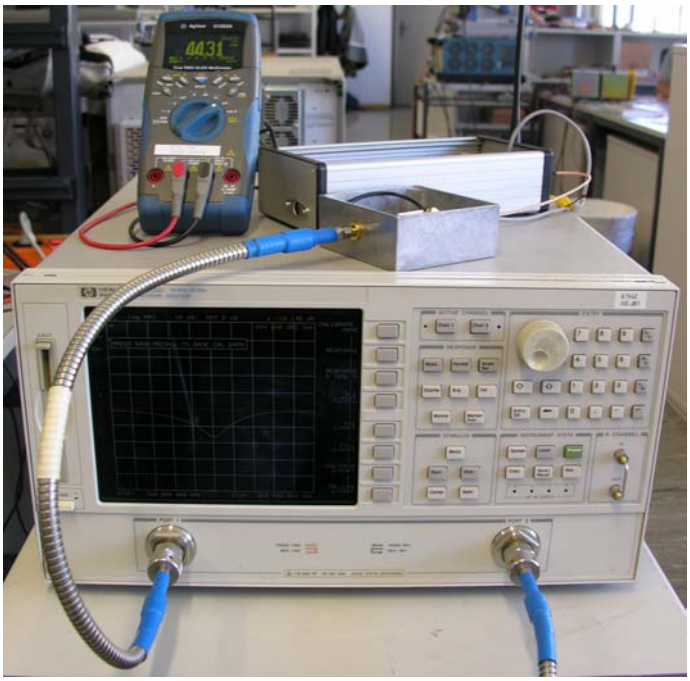

Figure 4.9: Experimental set up of cluster node to the Agilent 8720D Network Analyzer for evaluation of OOK demodulator return loss.

implication of this circuit behaviour is that the impedance matching network must be designed for a specific input power (e.g. −30dBm or less) and with a wide enough bandwidth to ensure a low return loss over a range of signal input powers.

The return loss of the power-optimised OOK demodulation circuit detailed in Section 3.2.2.5 was also measured using the identical test configuration. Figure 4.11 illustrates the return loss under the same frequency sweep and input power range as in Figure 4.10. The return loss profile is significantly improved compared to the original OOK demodulator circuit, where the magnitude of the return loss at the carrier frequency is between −10.9dB and −17.7dB over the range of signal input powers. This is a gain in return loss of up to +5dB, which results in an improvement in the receiver sensitivity of the wake-up radio when using the power-optimised OOK demodulator (as investigated in Section 4.2.2.2).

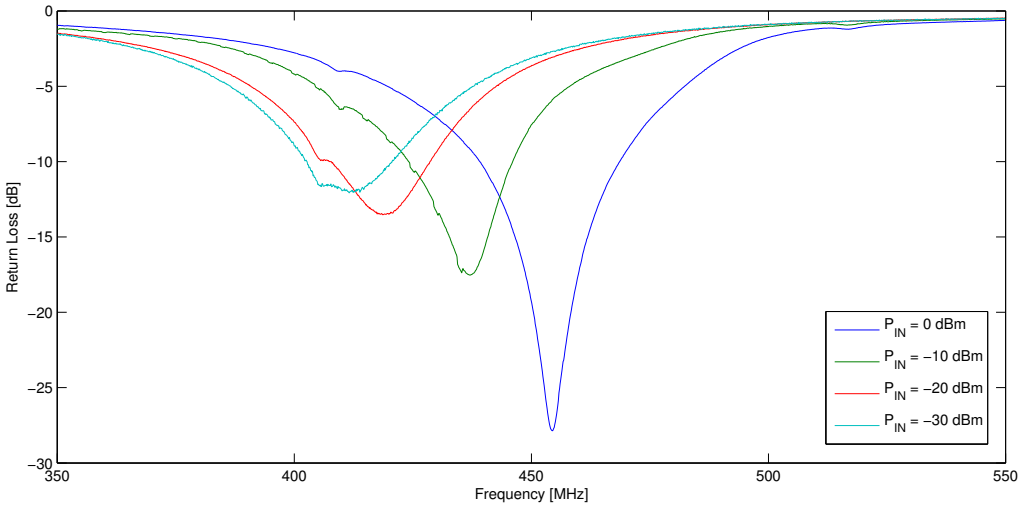

Figure 4.10: Return loss of the OOK demodulator circuit.

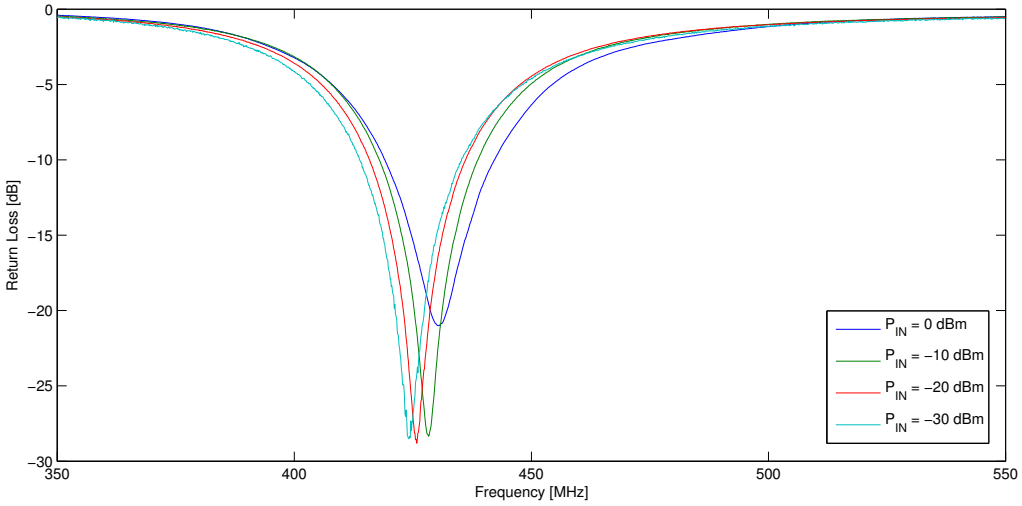

Figure 4.11: Return loss of the power-optimised OOK demodulator circuit.

4.2.2.2 Receiver Sensitivity

The range of the wake-up radio determines the operational range of the wireless hierarchical sensor network. As previously discussed in Section 2.1.3.4, the wake-up channel limits the range of the network due to the wake-up radio having a poor receiver sensitivity. Therefore, measuring the wake-up radio receiver sensitivity is a key performance metric for the wireless hierarchical sensor network.

The receiver sensitivity is defined as the minimum signal power at the receiver such that the receiver operates successfully. In the context of the wake-up radio, the receiver sensitivity is the minimum received signal power where a wake-up sequence is successfully received by the OOK demodulator and detected by the wake-up receiver. The wake-up radio receiver sensitivity is measured by connecting a variable attenuator between the cluster-head and the cluster node using coaxial cables, then for each increment of channel attenuation, a wake-up sequence is sent and the reception of measurement data is observed at the cluster-head. The receiver sensitivity is then equal to the power of the transmitted wake-up sequence minus the maximum channel attenuation where a successful wake-up sequence is observed. Coaxial cables are used to interconnect the variable attenuator so to remove any physical channel influences from the receiver sensitivity measurement. Figure 4.12 illustrates the experimental set up used to measure the wake-up radio sensitivity on the prototype wireless hierarchical sensor network.

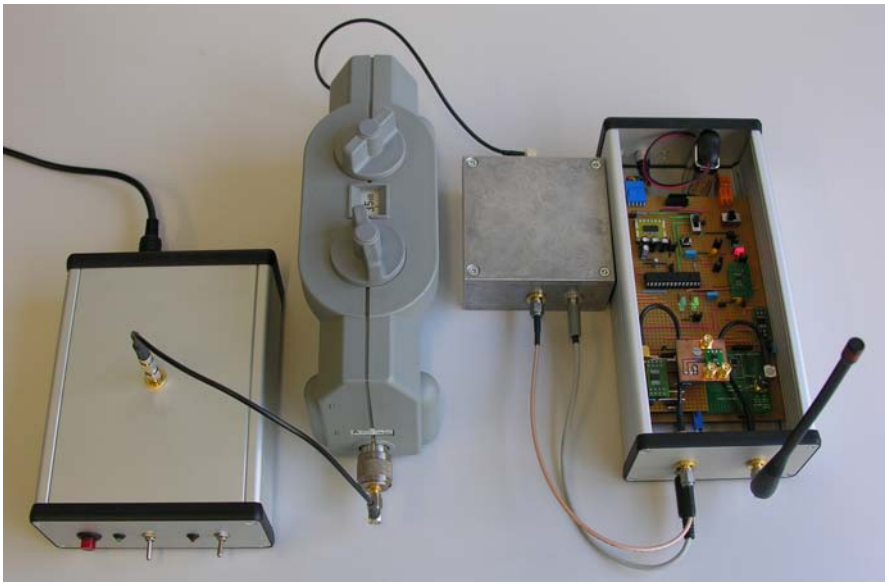

Figure 4.12: Experimental set up to measure the cluster node wake-up radio receiver sensitivity, consisting of (left-to-right) cluster-head, variable attenuator, OOK demodulator and cluster node.

The cluster-head was configured to transmit the wake-up sequence at the maximum power of +10dBm. Repeated experiments demonstrated that the wake-up radio received a transmitted wake-up sequence with +41dB attenuation, which is equivalent to a receiver sensitivity of −31dBm. The same experiment was performed with the power-optimised OOK demodulator circuit. Repeated experiments verified that a wake-up sequence was successfully received with up to +53dB attenuation, which is equivalent to a receiver sensitivity of −43dBm.

The experiments highlight a +12dBm improvement in receiver sensitivity when using the power-optimised OOK demodulator circuit.

Experiments in the laboratory indicate that the prototype wireless hierarchical sensor network using the power-optimised OOK demodulator and a quarter-wavelength monopole antenna with ground plane, has a communication range up to 14 meters line-of-sight.

Despite successful wireless communication, the wake-up radio receiver sensitivity is less than anticipated (i.e. less than the −52dBm published in [45]). This discrepancy is attributed to the impedance mismatch in the OOK demodulator PCB. The reduced communication range is related to the reduced receiver sensitivity and the sub-optimal antenna grounding on both the cluster node and cluster-head prototypes. It is anticipated that a multi-layer PCB with precision trace layout and adequate antenna grounding will significantly improve the range of the wireless hierarchical sensor network prototypes.

Chapter 5

Conclusion & Future Work

The new concept of a wireless hierarchical sensor network has been introduced in this thesis. The functional requirements, design, implementation and evaluation of a fully-functional prototype have been presented in the context of a ground surface temperature sensing use case.

The prototype hierarchical sensor network consists of an ultra-low power cluster node, and a stand-alone cluster-head. A wake-up radio architecture provides energy efficient wireless data delivery between the cluster node and cluster-head, with up to 14 meters line-of-sight communication range.

Ultra-low power system design principles have been successfully applied to the design and implementation of the cluster node. Assuming a deployment configuration of 1 measurement per hour with data requested once every 3 hours, the cluster node has an exceptionally low total average power dissipation of 78.5μW, leading to an estimated operation lifetime of more than 3 years from a 1000mAh coin cell battery. An optimistic power analysis of the cluster-head estimates greater than 3 years operation lifetime on a single battery charge, depending on the specific deployment configuration.

The cluster node achieves all functional requirements, with the exception of communication range. It is anticipated that a custom multi-layer PCB integrating analog and digital circuits with precision impedance matching and sufficient antenna grounding will significantly improve the wake-up radio receiver sensitivity, and hence extend the wake-up communication range up to 30 meters.

There are several directions for future work, covering both implementation-level optimisations to the existing prototypes and system-level aspects of wireless hierarchical sensor networking, which include:

- Investigate extending the wake-up channel into a generalised control channel supporting user-defined commands.
- Investigate the use of the wake-up radio architecture as a time synchronisation primitive in wireless sensor networks.

- Integrate the prototype circuitry onto a multi-layer PCB incorporating improved impedance matching and antenna grounding.
- Integration of the cluster-head prototype into an existing wireless sensor node suitable for field deployment.
- Design and manufacture mechanics for cluster node suitable for field deployment.
- Improve receiver sensitivity using a cascade voltage doubler circuit.
- Incorporate noise and interference rejection circuitry on the cluster node RF front-end.
- Increase data channel rate and incorporate a wake-up channel quality indicator.

The functional cluster node and cluster-head prototypes designed, developed and evaluated in this thesis successfully demonstrate the feasibility of wireless hierarchical sensing. The wake-up radio architecture enables greater flexibility in sensor placement, while the cluster node's ultra-low power autonomous operation improves sensing reliability compared to traditional mote-based wireless sensor networks.

Appendix A

Link Budget Analysis

A.1 Wake-up Channel

Link Characteristics

Carrier Frequency	434	MHz
Carrier Wavelength	0.691	m
Link Distance	30	m
Path Loss Exponent	2	
Date Rate	250	kbps
Noise-equivalent Bandwidth	30000	kHz
Temperature	80	°C

Transmitter

Transmission Power	10.0	dBm
Cable Losses	-0.5	dB
Transmission Antenna Gain	0.0	dB
Effective Isotropically Radiated Power	**9.5**	**dBm**

Channel

Free-space Path Loss	-54.7	dB
Multi-path Loss	-2.0	dB
Rain / Fog / Snow / Ice Losses	-2.0	dB
Total Path Loss	**-58.7**	**dB**

Receiver

Receiver Antenna Gain	0.0	dB
Polarization Loss	-0.2	dB
Cable Losses	-0.5	dB
Received Signal Level	**-49.9**	**dBm**
Receiver Sensitivity	**-52.0**	**dBm**
Link Margin	**2.1**	**dB**
Receiver Noise	-128.4	dBm
Receiver Noise Figure	15.0	dB
Minimum Signal-to-Noise Ratio	8.0	dB
Minimum Detectable Signal	**-105.4**	**dBm**

Figure A.1: Link budget analysis of the wake-up channel on the 434MHz ISM frequency band.

A.2 Data Channel

Link Characteristics

Carrier Frequency	434	MHz
Carrier Wavelength	0.691	m
Link Distance	30	m
Path Loss Exponent	2	
Date Rate	1.2	kbps
Noise-equivalent Bandwidth	1.9	kHz
Temperature	80	°C

Transmitter

Transmission Power	10.0	dBm
Cable Losses	-0.5	dB
Transmission Antenna Gain	0.0	dB
Effective Isotropically Radiated Power	**9.5**	**dBm**

Channel

Free-space Path Loss	-54.7	dB
Multi-path Loss	-2.0	dB
Rain / Fog / Snow / Ice Losses	-2.0	dB
Total Path Loss	**-58.7**	**dB**

Receiver

Receiver Antenna Gain	0.0	dB
Polarization Loss	-0.2	dB
Cable Losses	-0.5	dB
Received Signal Level	**-49.9**	**dBm**
Receiver Sensitivity	**-110.0**	**dBm**
Link Margin	**60.1**	**dB**
Receiver Noise	-170.3	dBm
Receiver Noise Figure	15.0	dB
Minimum Signal-to-Noise Ratio	8.0	dB
Minimum Detectable Signal	**-147.3**	**dBm**

Figure A.2: Link budget analysis of the data channel on the 434MHz ISM frequency band.

Appendix B

Design Space Exploration

B.1 Microcontroller Survey

Device	**AT32UC3L0128**	**MSP430F5418A**	**PIC24FJ128GA310**
Manufacturer	**Atmel**	**Texas Instruments**	**Microchip**
Architecture	32-bit	16-bit	16-bit
Max. System Clock	50 MHz	25 MHz	32 MHz
Flash	128 KB	128 KB	128 KB
SRAM	32 KB	16 KB	8 KB
Supply Voltage	1.6 – 3.6	1.8 – 3.6	2.0 – 3.6
Active	4.8 mA	1.8 mA	1.25 mA
Sleep	7.3 uA	2.1 uA	0.74 uA
	F=50MHz, Vcc=3.3V, w/ RTC	F=8MHz, Vcc=3.3V, w/ RTC	F=8MHz, Vcc=3.3V, w/ RTC
UART	4	2	4
I2C	2	2	2
SPI	1	2	2
External Interrupts	6	16	5
Timers	6 (16-bit)	3 (16-bit)	5 (16-bit)
Real-Time Clock (RTC)	Yes	Yes	Yes
ADC	8-channel @ 12-bit	14-channel @ 12-bit	24-channel @ 12-bit
Operating Temperature	-40 to +85	-40 to +85	-40 to +85
Cost per unit	28.3 CHF	14.9 CHF	9.15 CHF

Figure B.1: Short-list of three microcontrollers for prototype implementation.

B.2 Data Radio Survey

Device	Manufacturer	Operating Frequency	RX / TX Mode	Max Data Rate [kbps]	Supply Voltage [V]	Power Idle	Power RX Active [mA]	Power TX Active [mA]	IO Interface	Modulation Type	RX Sensitivity	Max TX Power [dBm]	Operating Temperature [°C]
AS3930	Austriamicrosystems	110 - 150 kHz	RX Only	4	2.4 - 3.6	400nA	0.0035 / 0.0053	-	SPI Interface	ASK/OOK	100uVrms	-	-40 to +85
MICRF113	Micrel	300 – 450 MHz	TX Only	10	1.8 – 3.6	-	-	12.5	2-wire Serial Interface	ASK/OOK	-	8.6	-40 to +85
MICRF219A	Micrel	300 – 450 MHz	RX Only	20	3.0 – 3.6	13 uA	6	-	2-wire Serial Interface	ASK/OOK	-110 dBm	-	-40 to +105
MICRF213	Micrel	300 – 450 MHz	TX Only	64	1.5 – 3.6	2.7 mA	-	11.2	2-wire Serial Interface	ASK/OOK/FSK	-	10	-40 to +125
MICRF211	Micrel	433.92MHz	RX Only	10	3.0 – 3.6	0.5 uA	6	-	2-wire Serial Interface	ASK/OOK	-110 dBm	10	-40 to +105
CC1000	Texas Instruments	300 – 1000 MHz	RX / TX	76	2.1 – 3.6	0.2 uA	7.4	26.7	3-wire Serial Interface	FSK	-109 dBm	10	-40 to +85
SX1213	Semtech	300 – 510 MHz	RX Only	32	2.1 – 3.6	65 uA	3	-	SPI Interface	FSK/OOK	-110 dBm	-	-40 to +85
SX1220	Semtech	290 – 1020 MHz	TX Only	32	1.8 – 3.7	0.9 mA	-	33	SPI Interface	(G)FSK/(G)MSK/OOK	-	10	-40 to +85
SX1240	Semtech	240 – 960 MHz	TX Only	10	1.8 – 3.7	0.5 uA	-	16.5	2-wire Serial Interface	FSK/OOK	-	10	-40 to +85
CC110L	Texas Instruments	300 – 928 MHz	RX / TX	250	1.8 – 3.6	165 uA	16	29.2	SPI Interface	(G)FSK/OOK	-112 dBm	10	-40 to +85
CC113L	Texas Instruments	300 – 928 MHz	RX Only	250	1.8 – 3.6	165 uA	16	-	SPI Interface	(G)FSK/OOK	-112 dBm	-	-40 to +85
CC115L	Texas Instruments	300 – 928 MHz	TX Only	250	1.8 – 3.6	165 uA	-	29.2	SPI Interface	(G)FSK/OOK	-	10	-40 to +85
TH72032	Melexis	300 – 450 MHz	TX Only	40	1.95 – 5.5	0.2 uA	-	10.6	2-wire Serial Interface	ASK	-	8	-40 to +125
TH71102	Melexis	280 – 510 MHz	RX Only	280	2.3 – 5.5	0.1 uA	8.2	-	2-wire Serial Interface	FSK/ASK	-110 dBm	-	-40 to +85
RFB3PL	HopeRF Electronic	390 – 440 MHz	RX Only	10	2.2 – 3.6	0.9 uA	3	-	2-wire Serial Interface	ASK/OOK	-108 dBm	-	-35 to +85
MAX7030	Maxim	315 – 433 MHz	RX / TX	66	2.1 – 3.6	0.8 uA	6.4	12.4	2-wire Serial Interface	ASK/OOK	-113 dBm	10	-40 to +125
MAX7060	Maxim	280 – 450 MHz	TX Only	50	2.1 – 3.6	0.4 uA	-	22	2-wire Serial Interface	FSK/ASK	-	10	-40 to +125
MAX7057	Maxim	300 – 450 MHz	RX Only	100	2.1 – 3.6	2.5 uA	5.8	-	2-wire Serial Interface	ASK	-114 dBm	10	-40 to +85
TDA5340	Infineon	300 – 960 MHz	RX / TX	40	3.0 – 3.6	0.9 uA	11	12.5	SPI Interface	FSK/ASK	-111 dBm	10	-40 to +110
TDA5220	Infineon	302 – 870 MHz	RX Only	20	3.0 – 3.6	60 uA	7.6	-	SPI Interface	FSK/ASK	-102 dBm	-	-40 to +105
TDA5150	Infineon	300 – 928 MHz	TX Only	50	1.9 – 3.6	0.4 uA	-	13	SPI Interface	(G)FSK/OOK	-	10	-40 to +85
ADF7020-1	Analog Devices	135 – 650 MHz	RX / TX	64	2.3 – 3.6	0.1 uA	17.6	21	SPI Interface	FSK/OOK	-112 dBm	10	-40 to +85

Figure B.2: Suvery of commercial non-volatile memory modules for prototype implementation.

Packet length [bits]	300
TX startup overhead [%]	20%
TX shutdown overhead [%]	20%
Fraction of max data rate [%]	50%

Device	Manufacturer	RX / TX Mode	Max Data Rate [kbps]	Power TX Active [mA]	Energy Metric [mA/packet]
MICRF113	Micrel	TX Only	10	12.5	1.050
MICRF112	Micrel	TX Only	50	11.2	0.188
SX1230	Semtech	TX Only	32	33	0.866
SX1240	Semtech	TX Only	10	16.5	1.386
CC110L	Texas Instruments	RX / TX	250	29.2	0.098
CC115L	Texas Instruments	TX Only	250	29.2	0.098
CC1000	Texas Instruments	RX / TX	76	26.7	0.295
TH72012	Melexis	TX Only	40	10.6	0.223
MAX7060	Maxim	TX Only	50	22	0.370
TDA5150	Infineon	TX Only	50	13	0.218
ADF7020-1	Analog Devices	RX / TX	64	21	0.276
TDA5340	Infineon	RX / TX	40	12.5	0.263
MAX7030	Maxim	RX / TX	66	12.4	0.158

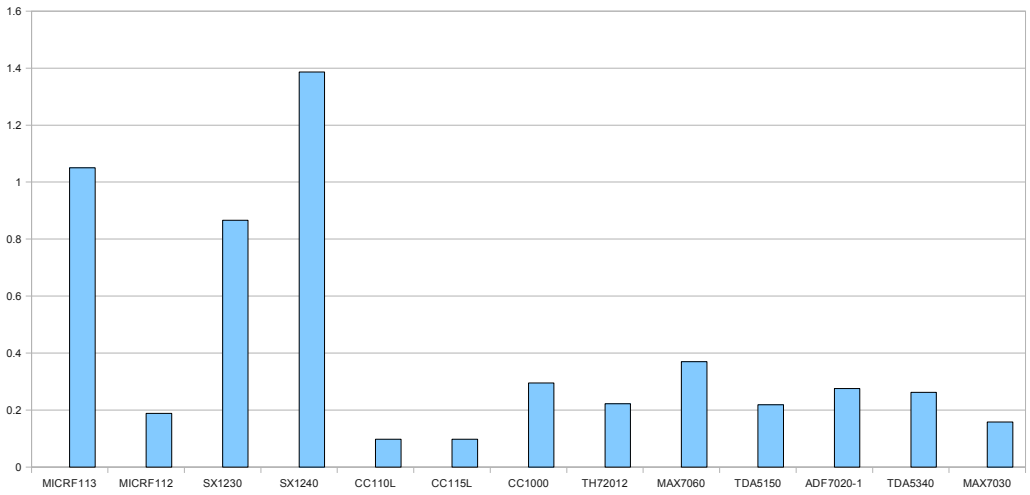

Figure B.3: Selection of most suitable data radio through minimisation of performance metric. The CC110L and CC115L modules from Texas Instruments are selected for prototype implementation.

B.3 Non-volatile Memory Survey

Device	Manufacturer	Technology	Size	I/O Interface	Supply Voltage [V]	Standby Current [uA]	Read Current [mA]	Write Current [mA]	Operating Temperature [°C]
SST25LF020A	Microchip	Serial Flash	2 Mb	SPI	3.0 – 3.6	15	10	30	-40 to +85
SST25VF010A	Microchip	Serial Flash	1 Mb	SPI	2.7 – 3.6	15	10	30	-40 to +85
SST25WF512	Microchip	Serial Flash	512 Kb	SPI	1.65 – 1.95	2	4	6	-40 to +85
AT24C1024B	Atmel	Serial EEPROM	1 Mb	I2C	1.8 – 3.6	4	2	3	-55 to +125
AT25512	Atmel	Serial EEPROM	512 Kb	SPI	1.8 – 5.5	2	5	5	-55 to +125
M24M01-R	ST	Serial EEPROM	1 Mb	I2C	1.8 – 5.5	5	1.5	2	-40 to +85
M95M01-R	ST	Serial EEPROM	1 Mb	SPI	1.8 – 5.5	5	5	5	-40 to +85
24AA1025	Microchip	Serial EEPROM	1 Mb	I2C	1.7 – 5.5	5	0.45	5	-40 to +125
24AA512	Microchip	Serial EEPROM	512 Kb	I2C	1.7 – 5.5	1	0.4	5	-40 to +125
FM25V10	Ramtron	Serial FRAM	1 Mb	SPI	2.0 – 3.6	5	0.3	0.3	-40 to +85

Figure B.4: Suvery of commercial non-volatile memory modules for prototype implementation.

Appendix C

Circuit Schematics

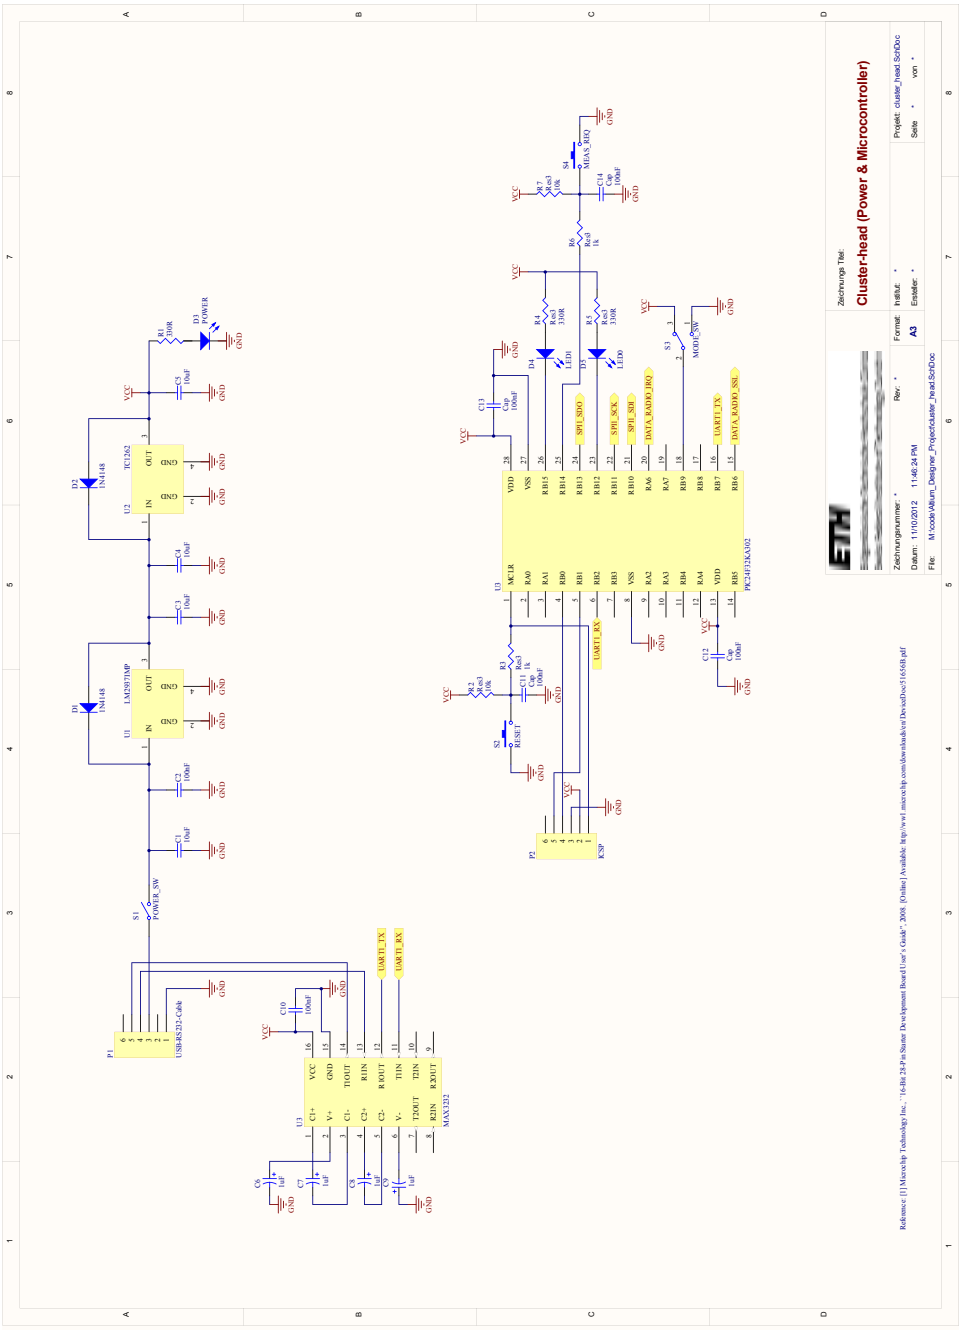

Figure C.1: Circuit schematic of cluster-head power supply and microcontroller.

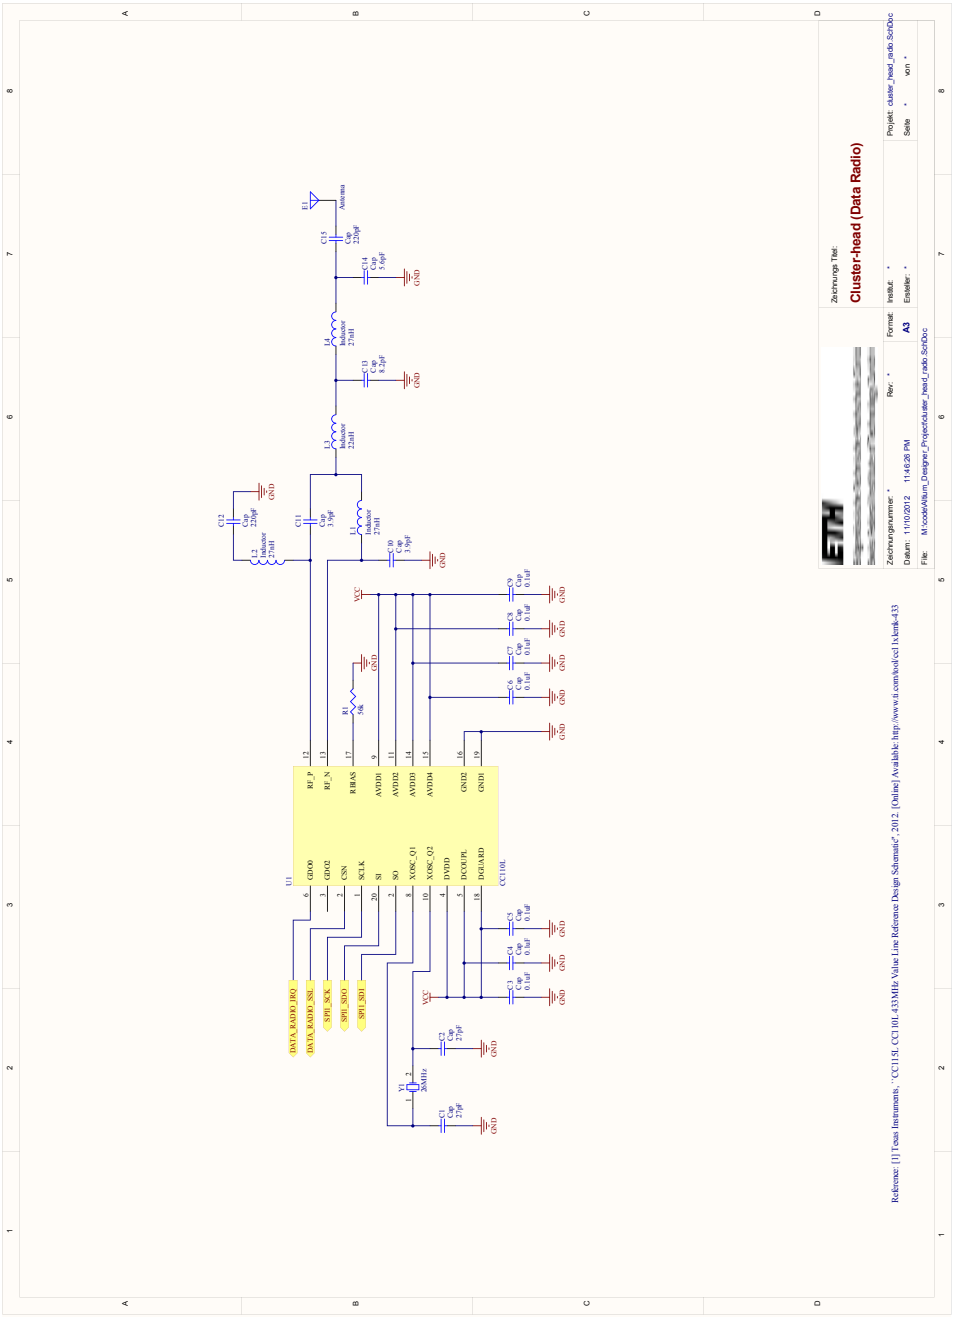

Figure C.2: Circuit schematic of cluster-head data radio.

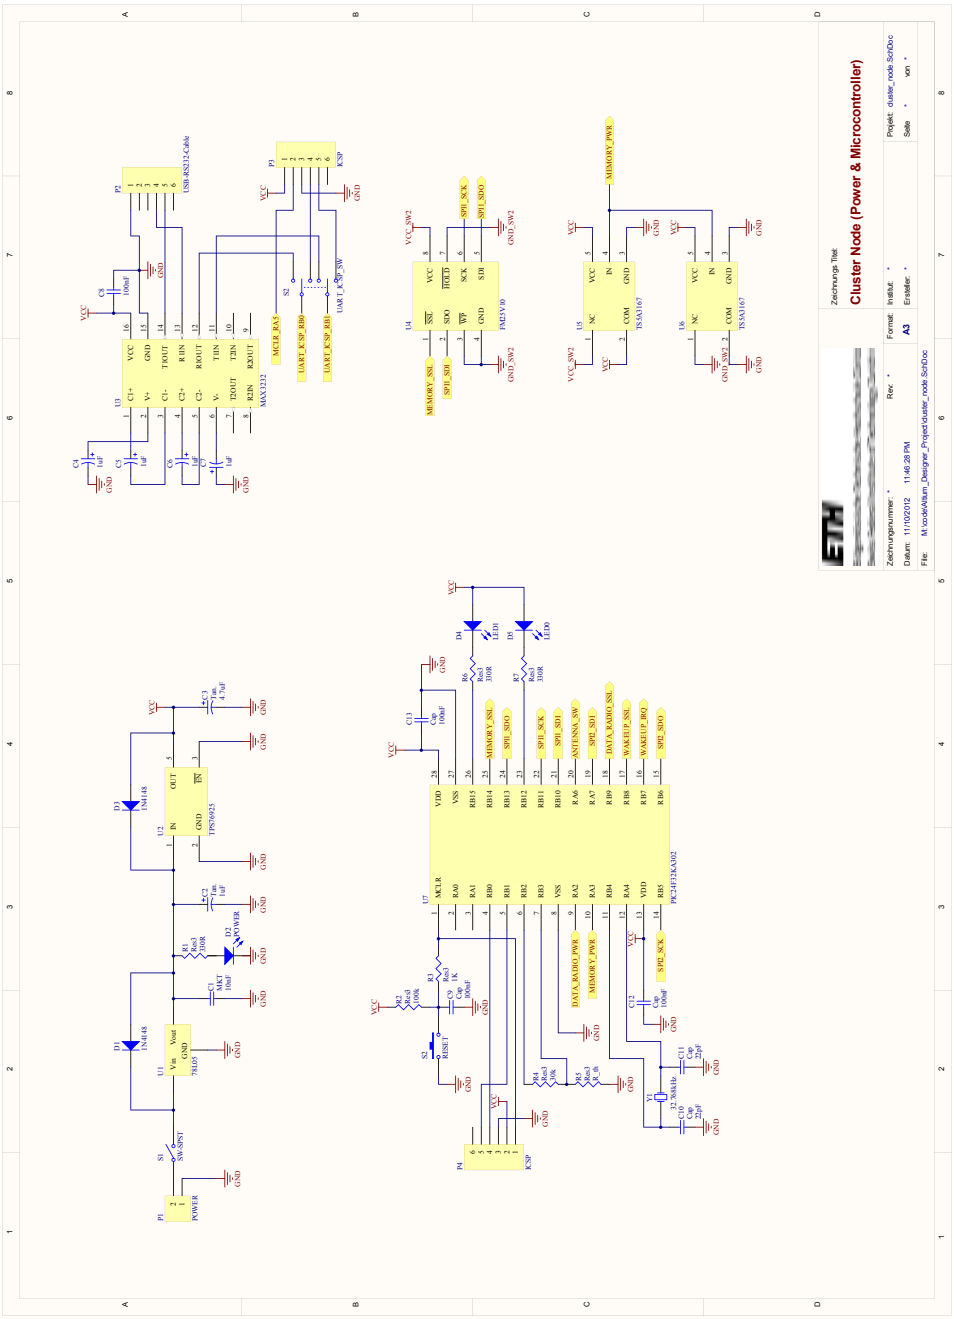

Figure C.3: Circuit schematic of cluster node power supply and microcontroller.

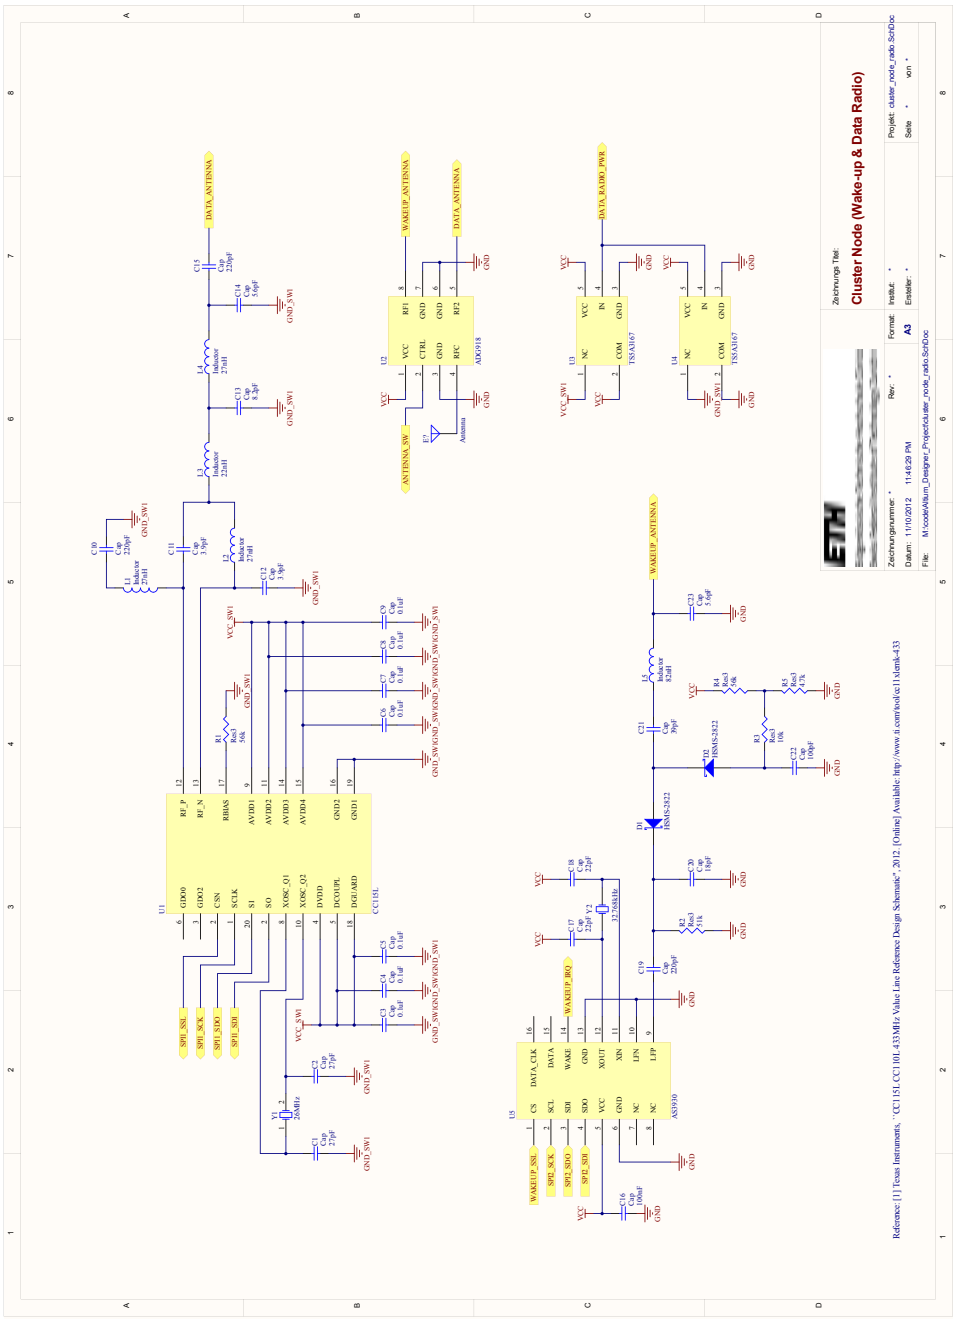

Figure C.4: Circuit schematic of cluster node wake-up radio (with biased diodes) and data radio.

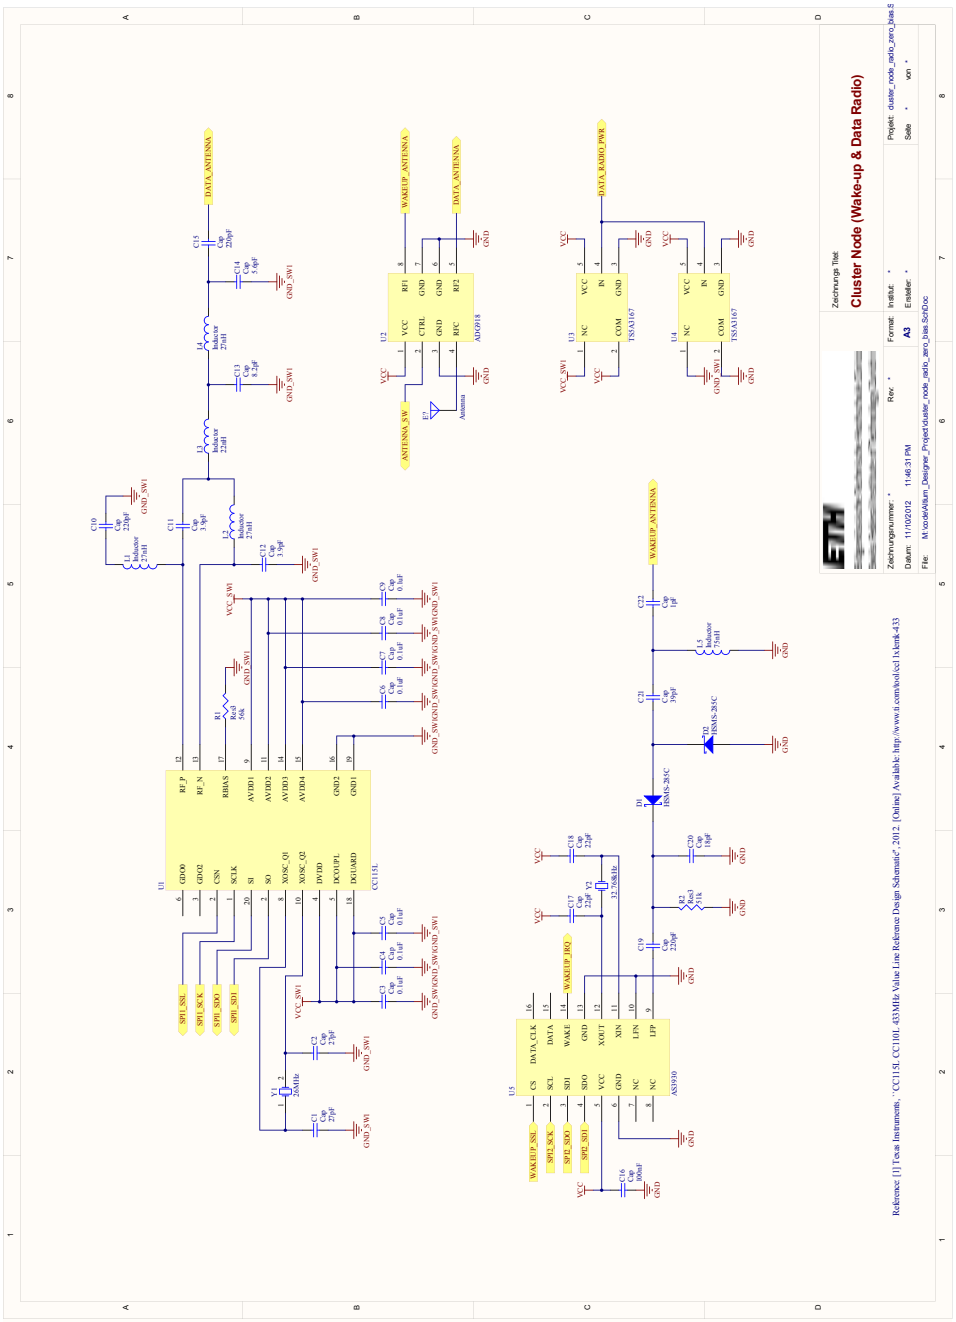

Figure C.5: Circuit schematic of cluster node wake-up radio (with zero biased diodes) and data radio.

Bibliography

[1] P. L. D. Abrie, *The Design of Impedance-Matching Networks for Radio-Frequency and Microwave Amplifiers*, Artech House Inc., 1985.

[2] E. J. Angelo, *Electronic Circuits: A Unified Treatment of Vacuum Tubes and Transistors*, McGraw-Hill, 1958.

[3] C. Bowick, J. Blyler, C. Ajluni, *RF Circuit Design*, Newnes, 2008.

[4] J. Carr, *RF Components and Circuits*, Newnes, 2002.

[5] J. Davis, *High-Speed Digital System Design*, Morgan & Claypool, 2006.

[6] K. Finkenzeller, *RFID Handbook: Fundamentals and Applications in Contactless Smart Cards, Radio Frequency Identification and Near-Field Communication*, John Wiley & Sons, Inc., 2010.

[7] J. Ganssle, *The Art of Designing Embedded Systems*, Newnes, 2000.

[8] P. Horowitz, W. Hill, *The Art of Electronics*, Cambridge, 1980.

[9] H. Karl, A. Willig, *Protocols and Architectures for Wireless Sensor Networks*, Wiley, 2007.

[10] P. Marwedel, *Embedded System Design*, Springer, 2006.

[11] K. McClaning, *Wireless Receiver Design for Digital Communications*, SciTech Publishing, 2012.

[12] E. McCune, *Practical Digital Wireless Signals*, Cambridge University Press, 2010.

[13] B. Otis, J. Rabaey, *Ultra-Low Power Wireless Technologies for Sensor Networks*, Springer, 2007.

[14] P. Z. Peebles, *Digital Communication Systems*, Prentice-Hall, 1987.

[15] G. Pottie, W. Kaiser, *Principles of Embedded Networked Systems Design*, Cambridge University Press, 2005.

[16] J. Rabaey, *Low Power Design Essentials*, Springer, 2009.

[17] C. W. Sayre, *Complete Wireless Design*, McGraw-Hill, 2008.

[18] J. Seybold, *Introduction to RF Propagation*, Wiley, 2005.

[19] D. E. Simon, *An Embedded Software Primer*, Addison-Wesley, 1999.

[20] R. N. Simons, *Coplanar Waveguide Circuits, Components, and Systems*, John Wiley & Sons, 2001.

[21] I. Sommerville, *Software Engineering (6th Edition)*, Addison-Wesley, 2001.

[22] D. Tse, P. Viswanath, *Fundamentals of Wireless Communication*, Cambridge University Press, 2005.

[23] B. C. Wadell, *Transmission Line Design Handbook*, Artech House, 1991.

[24] J. S. Wilson, *Sensor Technology Handbook*, Elsevier, 2005.

[25] P. Wilson, *The Circuit Designer's Companion (3rd Edition)*, Elsevier, 2012.

[26] W. Wolf, *Computers as Components: Principles of Embedded Computing System Design*, Elsevier, 2008.

[27] M. Yarvis, W. Ye, *Handbook of Sensor Networks: Compact Wireless and Wired Sensing Systems*, CRC Press, 2004.

[28] A. A. Abbasi, M. Younis, "A survey on clustering algorithms for wireless sensor networks", *Computer Communication, Vol. 30, No. 14-15, 2826-2841*, June 2007.

[29] I. Akyildiz, E. P. Stuntebeck, "Wireless Underground Sensor Networks: Research Challenges", *Ad Hoc Networks, Vol. 4, No. 6, Elsevier*, November 2006.

[30] J. Ansari, D. Pankin, P. Mähönen, "Radio-triggered Wake-ups with Addressing Capabilities for Extremely Low Power Sensor Network Applications", *International Journal of Wireless Information Networks, Vol. 16, No. 3*, 2009.

[31] J. Beutel *et al.*, "PermaDAQ: A Scientific Instrument for Precision Sensing and Data Recovery in Environmental Extremes", *International Conference on Information Processing in Sensor Networks*, April 2009.

[32] J. Beutel, B. Buchli, F. Ferrari, M. Keller, M. Zimmerling, L. Thiele, "X-SENSE: Sensing in Extreme Environments", *Proceedings of Design, Automation and Test in Europe*, March 2011.

[33] E. Björnemo, "Energy Constrained Wireless Sensor Networks: Communication Principles and Sensing Aspects", *PhD Thesis*, January 2009.

[34] B. Buchli, F. Sutton, J. Beutel, "GPS-Equipped Wireless Sensor Network Node for High-Accuracy Positioning Applications", *European Conference on Wireless Sensor Networks*, February 2012.

[35] N. Burri, P. von Rickenbach, R. Wattenhofer, "Dozer: Ultra-Low Power Data Gathering in Sensor Networks", *Proceedings of the 6th International Conference on Information Processing in Sensor Networks*, April 2007.

[36] A. Caracas *et al.*, "Energy-Efficiency Through Micro-Managing Communication and Optimizing Sleep", *IEEE Conference on Sensor, Mesh and Ad Hoc Communications and Networks*, June 2011.

[37] J. L. da Silva Jr. *et al.*, "Design Methodology for PicoRadio Networks", *Proceedings of the Design Automation and Test in Europe*, March 2001.

[38] I. Demirkol, C. Ersoy, F. Alagoz, "MAC Protocols for Wireless Sensor Networks: A Survey", *IEEE Communications Magazine, Vol. 44, Issue 4, 115 - 121*, April 2006.

[39] I. Demirkol, C. Ersoy, E. Onur, "Wake-up Receivers for Wireless Sensor Networks: Benefits and Challenges", *IEEE Wireless Communications*, August 2009.

[40] W. Dong, C. Chen, X. Liu, J. Bu, "Providing OS Support for Wireless Sensor Networks: Challenges and Approaches", *IEEE Communications Surveys & Tutorials, Vol. 12, No. 4*, May 2010.

[41] M. S. Durante, S. Mahlknecht, "An ultra low power Wakeup Receiver for Wireless Sensor Nodes", *Third International Conference on Sensor Technologies and Applications*, June 2009.

[42] S. Evans, B. M. E. Smith, "A Radio Echo Equipment for Depth Sounding in Polar Ice Sheets", *Journal of Scientific Instruments, Series 2, Vol. 2*, 1969.

[43] G. Frantz, J. Henkel, J. Rabaey, T. Schneider, M. Wolf, U. Batur, "Ultra-Low Power Signal Processing", *IEEE Signal Processing Magazine, 149-159*, March 2010.

[44] B. Gadek, J. Leszkiewicz, "Infulence of snow cover on ground surface temperature in the zone of sportadic permafrost, Tatra Mountains, Polank and Slovakia", *Cold Regions Science and Technology, 60, 205-211*, 2010.

[45] G. U. Gamm, M. Sippel, M. Kostic, L. M. Reindl, "Low-power sensor node with addressable wake-up on-demand capability", *International Journal of Sensor Networks, Vol. 11, No. 1*, January 2012.

[46] S. Gubler, J. Fiddes, M. Keller, S. Gruber, "Scale-dependent measurement and analysis of ground surface temperature variability in alpine terrain", *The Cryosphere, 5, 431-443*, 2011.

[47] D. Harel, A. Pnueli, "On the Development of Reactive Systems", *Logics and Models of Concurrent Systems, Vol. F13, 477 - 498, Springer-Verlag Berlin Heidelberg*, 1985.

[48] J. K. Hart *et al.*, "A Wireless Multi-sensor Subglacial Probe: Design and Preliminary Results", *Journal of Glaciology, Vol. 52, No. 178*, 2006.

[49] J. K. Hart, K. Martinez, "Environmental Sensor Networks: A Revolution in the Earth System Science?", *Earth-Science Reviews, Vol. 78, 177-191*, October 2006.

[50] C. Hambeck, S. Mahlknecht, T. Herndl, "A 2.4uW Wake-up Receiver for Wireless Sensor Nodes with -71dBm Sensitivity", *IEEE International Symposium on Circuits and Systems*, May 2011.

[51] D. Hasenfratz, A. Meier, M. Woehrle, M. Zimmerling, L. Thiele, "Poster Abstract: If You Have Time, Save Energy with Pull", *Proceedings of the 8th ACM Conference on Embedded Networked Sensor Systems, 423-424*, November 2010.

[52] J. Hill, R. Szewczyk, A. Woo, S. Hollar, D. Culler, K. Pister, "System Architecture Directions for Networked Sensors", *Proceedings of the 9th International Conference on Architectural Support for Programming Languages and Operating Systems, 93-104*, November 2000.

[53] V. Jelicic, "Power Management in Wireless Sensor Networks with High-Consuming Sensors", *Qualifying Doctoral Examination, University of Zagreb*, April 2011.

[54] P. Le-Huy, S. Roy, "Low-Power 2.4GHz Wake-up Radio for Wireless Sensor Networks", *IEEE International Conference on Wireless and Mobile Computing*, October 2008.

[55] A. Lewkowicz, "Evaluation of Miniature Temperature-loggers to Monitor Snowpack Evolution at Mountain Permafrost Sites, Northwestern Canada", *Permafrost and Periglacial Process, 19, 323-331*, 2008.

[56] S. Liang, Y. Tang, Q. Zhu, "Passive Wake-up Scheme for Wireless Sensor Networks", *Second International Conference on Innovative Computing, Information and Control*, September 2007.

[57] M. Lont *et al.*, "Ultra-Low Power FSK Wake-up Receiver Front-End for Body Area Networks", *IEEE Radio Frequency Integrated Circuits Symposium*, June 2011.

[58] S. Marinkovic, E. M. Popovici, "Nano-Power Wireless Wake-Up Receiver with Serial Peripheral Interface", *IEEE Journal of Selected Areas in Communications, Vol. 29, No. 8*, September 2011.

[59] J.-C. Otto *et al.*, "Detection of Mountain Permafrost by Combining High Resolution Surface and Subsurface Information - An Example from The Glatzbach Catchment, Austrian Alps", *Geografiska Annaler: Series A, Physical Geography, 94, 43-57*, 2012.

[60] J. Richmond, J. Rabaey, "Digital Energy Detection for OOK Demodulation in Ultra-Low Power Radios", *IEEE International Symposium on Circuits and Systems*, May 2011.

[61] M.-O. Schmid, S. Gubler, J. Fiddes, S. Gruber, "Inferring snow pack ripening and melt out from distributed ground surface temperature measurements", *The Cryosphere Discussions, 6, 563-591*, 2012.

[62] C. E. Shannon, "A Mathematical Theory of Communication", *Bell System Technical Journal, Vol. 27, 379-423 and 623-656*, July and October 1948.

[63] W.-C. Shih, R. Jurdak, B.-H. Lee, D. Abbott, "High sensitivity wake-up radio using spreading codes: design, evaluation , and applications", *EURASIP Journal on Wireless Communications and Networking*, 2011.

[64] F. Sutton, "Wireless Torque Measurement System", *Honours Engineering Project, The University of Adelaide*, September 2004.

[65] F. Sutton, "GPS-equipped Wireless Sensor Node for PermaSense", *ETH Semester Thesis*, December 2010.

[66] F. Sutton, "Energy Harvesting Wireless Anemometer", *ETH Semester Thesis*, July 2011.

[67] B. van der Doorn, "A prototype low-cost wakeup radio for the 868 MHz band", *International Journal of Sensor Networks, Vol. 6, No. 1*, February 2009.

[68] J. Yick, B. Mukherjee, D. Ghosal, "Wireless Sensor Network Survey", *Computer Networks Vol. 52, 2292–2330*, April 2008.

[69] D. McMahill, "Coplanar Waveguide Analysis/Synthesis Calculator", 2009. [Online] Available: `http://wcalc.sourceforge.net/cgi-bin/coplanar.cgi`

[70] Geotest AG, "UTL-3 Scientific Datalogger", `http://www.utl.ch/`

[71] Proxxon GmbH, "Proxxon Micro Miller MF70", `http://www.proxxon.com/eng/html/27110.php`

[72] X-Sense Project, http://www.permasense.ch/

[73] Analog Devices, Inc., "ADG918/ADG919 Datasheet", 2008. [Online] Available: http://www.analog.com/static/imported-files/data_sheets/ADG918_919.pdf

[74] Antenna Factor, "433MHz 1/4-Wavelength Monopole Antenna (ANT-433-CW-HW) Datasheet", 2010. [Online] Available: http://www.linxtechnologies.com/resources/data-guides/ant-433-cw-hw.pdf

[75] Atmel Corporation, "ATA5283 Interface IC for 125kHz Wake-up Function", 2004. [Online] Available: http://www.datasheetcatalog.org/datasheet/atmel/doc4598.pdf

[76] austriamicrosystems AG, "AS3930 Single Channel Low Frequency Wakeup Receiver", 2009. [Online] Available: http://www.ams.com/eng/content/download/23692/414425/17914

[77] Avago Technologies, "HSMS-282x Surface Mount RF Schottky Barrier Diodes", 2009. [Online] Available: http://www.avagotech.com/docs/AV02-1320EN

[78] Avago Technologies, "HSMS-285x Series Surface Mount Zero Bias Schottky Detector Diodes", 2009. [Online] Available: http://www.avagotech.com/docs/AV02-1377EN

[79] Avago Technologies, "Linear Models for Diode Surface Mount Packages - Application Note 1124", 2010. [Online] Available: http://www.avagotech.com/docs/AV02-0038EN

[80] CEPT, "ERC Recommendation 70-03", October 2012. [Online] Available: http://www.erodocdb.dk/docs/doc98/official/pdf/rec7003e.pdf

[81] ETSI, "ETSI EN 300 220-1", January 2012. [Online] Available: http://www.etsi.org/deliver/etsi_en/300200_300299/30022001/02.04.01_40/en_30022001v020401o.pdf

[82] ETSI, "ETSI EN 300 328", April 2012. [Online] Available: http://www.etsi.org/deliver/etsi_en/300300_300399/300328/01.08.01_30/en_300328v010801v.pdf

[83] FTDI Ltd., "USB to RS232 Serial Converter Range of Cables", 2010. [Online] Available: http://www.ftdichip.com/Support/Documents/DataSheets/Cables/DS_USB_RS232_CABLES.pdf

[84] Maxim Integrated, "DS1922L/DS1922T Temperature Logger iButton with 8KB Data-Log Memory", 2012. [Online] Available: http://datasheets.maximintegrated.com/en/ds/DS1922L-DS1922T.pdf

[85] Measurement Specialties, Inc., "44006RC Precision Epoxy NTC Thermistor", 2008. [Online] Available: http://www.meas-spec.com/WorkArea/showcontent.aspx?id=1690

[86] Microchip Technology Inc., "MPLAB C30 C Compiler User's Guide", 2007. [Online] Available: http://ww1.microchip.com/downloads/en/devicedoc/c30_users_guide_51284f.pdf

[87] Microchip Technology Inc., "16-Bit 28-Pin Starter Development Board User's Guide", 2008. [Online] Available: http://ww1.microchip.com/downloads/en/DeviceDoc/51656B.pdf

[88] Microchip Technology Inc., "MPLAB IDE User's Guide with MPLAB Editor and MPLAB SIM Simulator", 2009. [Online] Available: http://ww1.microchip.com/downloads/en/DeviceDoc/MPLAB_User_Guide_51519c.pdf

[89] Microchip Technology Inc., "PICkit 3 Programmer/Debugger User's Guide", 2010. [Online] Available: http://ww1.microchip.com/downloads/en/DeviceDoc/51795B.pdf

[90] Microchip Technology Inc., "PIC24FV32KA304 16-Bit Flash Microcontroller with XLP Technology Datasheet", 2011. [Online] Available: http://ww1.microchip.com/downloads/en/DeviceDoc/39995_b.pdf

[91] Onset Computer Corporation, "HOBO Water Temperature Pro v2 Data Logger", 2012. [Online] Available: http://www.onsetcomp.com/files/manual_pdfs/10366-F-MAN-U22-001.pdf

[92] Onset Computer Corporation, "HOBOnode Outdoor Wireless Systems", 2012. [Online] Available: http://www.onsetcomp.com/files/manual_pdfs/12473-K-MAN-HOBOnode-UG.pdf

[93] Ramtron Inc., "FM25V10 1Mb Serial 3V F-RAM Memory Datasheet", August 2008. [Online] Available: http://www.ramtron.com/files/pdfs/FM25V10ds_r1-0.pdf

[94] Sony Corporation, "CR2477 Lithium Manganese Dioxide Battery Datasheet", 2012. [Online] Available: http://www.sony.net/Products/MicroBattery/cr/pdf/cr2477.pdf

[95] Texas Instruments, "TPS76925 2.5V Ultra-low Power 100-mA Low Drop-out Linear Regulator", 2001. [Online] Available: http://www.ti.com/lit/gpn/tps76925

[96] Texas Instruments, "CC110L Value Line Transceiver", 2012. [Online] Available: http://www.ti.com/lit/ds/symlink/cc110l.pdf

[97] Texas Instruments, "CC115L Value Line Transmitter", 2012. [Online] Available: http://www.ti.com/lit/ds/symlink/cc115l.pdf

[98] Texas Instruments, "0.9-Ω SPST Analog Switch", 2012. [Online] Available: http://www.ti.com/lit/gpn/ts5a3167

[99] Texas Instruments, "Value Line Evaluation 433MHz Modules (CC11XLEMK-433)", 2012. [Online] Available: http://www.ti.com/tool/cc11xlemk-433