
D-1.1 – Report on Environmental and Platform Models

Grant Agreement no: 317826
www.relyonit.eu

Date: June 5, 2013

Author(s) and affiliation: Marco Antonio Zúñiga (TUD), Carlo Alberto Boano (UzL), James Brown (ULANC), Chamath Keppitiyagama (SICS), Felix Jonathan Oppermann (UzL), Paul Alcock (ULANC), Nicolas Tsiftes (SICS), Utz Roedig (ULANC), Kay Römer (UzL), Thimo Voigt (SICS), and Koen Langendoen (TUD).

Work package/task: WP1

Document status: Final

Dissemination level: Public

Keywords: Wireless sensor networks, temperature, interference, received signal strength, clock drift, analytical models, experiments.

Abstract This deliverable describes the progress made on the platform and environmental models during the first eight months of work. This progress is related to Tasks 1.1 and 1.2 in the Description of Work. The overarching goals of these two tasks are: (i) to capture the relevant temporal and spatial distribution of environmental properties, in particular, temperature and interference; and (ii) to develop platform models that capture how those environmental properties affect the signal strength, timing, sensing accuracy, and energy consumption of WSN hardware platforms. During these first eight months, our focus has been on analysing the effect of **temperature**. We now have a thorough understanding of the effect the temperature on the **signal strength** and a good understanding of the effect of temperature on **timing**. We have also started to evaluate the effect of temperature on sensing and energy consumption. With respect to **interference**, we have started to define some initial analytical models to capture the bursty nature of various interference sources.

Disclaimer

The information in this document is proprietary to the following RELYonIT consortium members: University of Lübeck, Swedish Institute of Computer Science AB, Technische Universiteit Delft, University of Lancaster, Worldsensing, Acciona Infraestructuras S.A.

The information in this document is provided “as is”, and no guarantee or warranty is given that the information is fit for any particular purpose. The user uses the information at his sole risk and liability. The above referenced consortium members shall have no liability for damages of any kind including without limitation direct, special, indirect, or consequential damages that may result from the use of these materials subject to any liability which is mandatory due to applicable law.

Copyright 2013 by University of Lübeck, Swedish Institute of Computer Science AB, Technische Universiteit Delft, University of Lancaster, Worldsensing, Acciona Infraestructuras S.A.

Contents

1	Introduction	8
1.1	Our progress thus far	9
2	Temperature - Platform Models	11
2.1	Impact of temperature on signal strength	11
2.1.1	Long-Term Outdoor Deployment	12
2.1.2	Controlled Testbed Experiments	14
2.1.3	Generalized model of the effect of temperature on RSS	19
2.2	Impact of Temperature on Radio Range Shrinkage and Network Topology ..	20
2.2.1	Radio Range Shrinkage and Loss of Links Due to Temperature	20
2.2.2	Impact of Temperature on Network Topology	24
2.3	Impact of temperature on clock drift	26
2.3.1	Experimental Results	27
	Primary System Clock	27
	Real-Time Clock	30
	Other Platform Clocks	31
2.3.2	Clock Drift Models	31
2.4	Impact of temperature on sensing	34
2.4.1	Experimental Results	35
2.4.2	Findings	36
2.5	Impact of temperature on energy supply and consumption	36
2.5.1	Experimental Results	37
2.5.2	Findings	39
3	Temperature - Environmental Models	40
3.1	Probabilistic Analysis	40
3.2	Network Profile	41
3.3	Node Profile	42
3.4	Aggregating models	44
4	Radio Interference	46
4.1	Platform Model	46
4.2	Environmental Model	49
4.2.1	Channel occupancy model	49
4.2.2	Node Models	51
4.2.3	Network Model	54
5	Conclusions	56

List of Figures

2.1	Temperature has a strong impact on the quality of links in our outdoor WSN. During daytime, when temperature is high, there is a significant reduction in PRR (a). Also the trend of RSSI and noise floor resembles the one of temperature, with a sharp decrease when temperature increases (b). Values are averaged over a timespan of 10 minutes.	13
2.2	The relationship between RSSI and temperature (a) and between noise floor and temperature (b) can be approximated as a linear function, and the trend is similar for different nodes.	13
2.3	Experimental setup in controlled testbed experiments.	15
2.4	Impact of temperature on the quality of links in our controlled testbed. We heat transmitter and receiver nodes separately first, and then both of them at the same time. When temperature increases, PRR, LQI, and RSSI decrease significantly, with the highest impact occurring when both nodes are heated at the same time. The periodic noise is due to a Wi-Fi access point beaconing in proximity of the testbed.	16
2.5	Figure (a) shows that the relationship between RSSI and temperature is similar when using different hardware platform and can be approximated as a linear function, but with different parameters. Figure (b) shows the non-linearities in the response of the CC2420 radio measured using Maxfor nodes. Temperature on the x-axis is computed as the average temperature of the transmitter and receiver temperature.	17
2.6	Relationship between RSSI, noise floor, SNR and temperature when transmitter (blue) and receiver (black) nodes are heated individually, and when both nodes (red) are heated at the same time.	18
2.7	The range shrinks with the increase of temperature.	21
2.8	A node loses links with the increase of temperature.	23
2.9	The network topology changes considerably within 11 hours.	25
2.10	The number of usable links goes down with the increase in temperature.	25
2.11	Maxfor MTM-CM5000MSP Primary System Clock	28
2.12	Zoltertia Z1 Primary System Clock	29
2.13	ST MB950 Primary System Clock	29
2.14	Maxfor MTM-CM5000MSP RTC Clock	30
2.15	CC2420 Radio Clock Period	31
2.16	Model Fitting for RC Clocks	32
2.17	Example of two different model that describe Frequency change with Temperature for RT Crystals of different cuts. These figures are found in [21]	33
2.18	Model Fitting for RT-Clocks	34
2.19	Maxfor RTC model Residuals	34
2.20	Maxfor ADC Sampling with Varying Temperature	35
2.21	Maxfor Energy Consumption vs Temperature	38

2.22	Zolertia Z1 Energy Consumption vs Temperature	39
2.23	ST MB950 Energy Consumption vs Temperature	39
3.1	Temperature of 16 nodes during a day in an outdoor deployment. A temperature of 0°C indicates that no data was collected during that time.	41
3.2	Periodic temperature pattern of a transmitter-receiver pair	43
4.1	Co-channel rejection as specified in the Texas Instruments CC2420 radio datasheet [41].	47
4.2	Co-channel rejection threshold found experimentally on the CC2420 transceiver, using Maxfor MTM-CM5000MSP nodes.	48
4.3	Example of interfering signal of strength R_{max} as recorded by a Maxfor MTM-CM5000MSP node performing a continuous RSSI noise floor measurements at a rate of 45.5 KHz. The shaded area shows the busy period of interference according to the two-state semi-Markov channel occupancy model. R_{THR} is set to -56 dBm.	50
4.4	RSSI values measured using off-the-shelf wireless sensor nodes operating in the 2.4 GHz ISM band. Please notice the different scale of the x -axis [11].	51
4.5	Worse-case scenario analysis for the transmission of a packet using CSMA-CA when using an interference model based on the minimum duration of an idle period min_{idle} and maximum duration of a busy period max_{busy}	52
4.6	Cumulative distribution function (CDF) of idle and busy periods measured by a Maxfor MTM-CM5000MSP node in the presence of a laptop continuously downloading a file from a nearby access point [11].	53
4.7	Cumulative distribution function (CDF) of busy periods among several Maxfor MTM-CM5000MSP nodes located inside a 56 m ² office. The colors identify the nodes in which similar interference patterns were measured.	55

List of Tables

1.1	Environmental Models	9
1.2	Platform Models	9
2.1	Largest temperature variations on a single node as seen in our outdoor deployment.	14
2.2	Determined RC temperature coefficients for the evaluated Maxfor MTM-CM5000MSP and Zolertia Z1 nodes.	32
2.3	Determined RT temperature coefficients for the evaluated Maxfor MTM-CM5000MSP nodes.	35

Executive Summary

The performance of wireless sensor and actuator networks (WSAN) depends on two fundamental factors: the hardware platform and the characteristic of the surrounding environment. A network of sensor nodes that is moved from an indoor commercial building to an outdoor forest will have a dramatically different performance. The role of the environment is central to the operation of these networks. But not only the environment plays a significant role, the hardware platform plays a major role too. If for a given environment, say an outdoor forest, we change the sensor nodes for a platform with a different radio transceiver or microprocessor, the performance can change dramatically as well. Our goal is to analyze, model and predict the effects of hardware and environmental factors on the performance of WSAN.

This report focuses on an environmental variable that has received little attention in the research community, but that nevertheless has a significant effect on the operation of the network: temperature. There is an extensive body of knowledge on the impact of several environmental variables on the performance of WSAN. By and large, the focus of these studies has been on analysing the degradation of link qualities caused by effects such as obstacles, foliage and climatological events such as fog and rain. Similarly to these environmental variables, temperature also affects link quality by reducing the received signal strength (RSS) by several dBs. But the impact of temperature does not stop there. Due to the general effect of temperature on electric conductors and semiconductors, changes in temperature can also affect the timing accuracy (clocks), sensing accuracy (analog to digital converters), and energy consumption (battery capacity) of the hardware platforms. These effects play a significant role on the performance of network protocols, on the accuracy of the delivered data, and on the lifetime of the network.

Our most important contribution is a thorough analysis of the effect of temperature on the received signal strength (RSS) and radio coverage of low-power radio transceivers. Our work is motivated by some initial studies reporting how temperature affects the quality of low-power links. These studies, however, were not systematic and the findings could not be generalized – the experiments were performed on a single platform and on a relatively short range of temperatures. To gain a deeper insight on the effect of temperature, we developed a low-cost yet accurate testbed that allows us to test a wide range of temperatures with fine granularity, from 0°C to 70°C. This testbed allows us to vary the on-board temperature of sensor nodes in a repeatable fashion, and we study systematically the impact of temperature on various sensor network platforms.

We show that temperature affects transmitting and receiving nodes differently, and that all platforms follow a similar trend that can be captured by a simple first-order model. This general model allows us to predict the RSS decay for a given temperature profile. For example, if a WSN network is deployed in an environment where the maximum and minimum day temperature lead to a delta of 40°C (a normal change considering that nodes exposed to sun light can reach high temperatures), our model can predict the change in RSS (for instance, a change of 5-6 dB for the Maxfor platform). Our contribution goes beyond the analysis of the RSS decay. We integrate our first-order model with the well-known log-normal path loss model to quantify the impact of temperature on the transmission range, i.e., the distance covered by a link. Quantifying the effect on the transmission range is central to understanding the effect on upper layer protocols, because different transmission ranges lead to fundamentally different topologies. Based on data from a year-long deployment in Sweden, we show how the topology of the network changes in a real scenario.

1 Introduction

This deliverable describes the progress made so far in modelling environmental phenomena and hardware platforms. While there is a long list of environmental phenomena affecting the performance of wireless sensor and actuator networks (WSAN), in RELYonIT we focus on two phenomena:

Temperature Since the inception of modern electronic devices, temperature has been known to affect their performance. The clearest and simplest example is the operational range given on the data sheets and user guides of all electronic devices. However, while the effect of temperature is well understood at low hardware levels (individual electrical and electronic components), the impact of these low-level effects on the overall networking performance of complex sensor systems has not been fully investigated.

Interference In the last decade, the widespread adoption of wireless technologies has increased the demand for bandwidth. This trend is leading to severe radio interference problems, especially, on the license-free ISM bands – a band that is used by most newly developed systems. Within the interference domain, our goal is to bridge the gap between (i) theoretical studies that rely on some unrealistic assumptions, and (ii) empirical studies that focus on single platforms and environments. We want to develop simple models that capture the impact of interference in a generic and platform-independent manner.

With regards to hardware platforms, our modelling efforts focus on four *pillars* that are fundamental for the operation of WSAN:

Communication Arguably the most fundamental capability required by a network is that of communicating with peers. In any computer network it is henceforth central to identify the elements that affect communication. In RELYonIT, we focus on modelling the effects that temperature and interference have on the received signal strength (RSS) and signal to noise ratio (SNR) of radio communication.

Timing Several protocols on the Data Link, Network and Transport Layers rely on time synchronization. Clock drift can affect the timely operation of the network. Our efforts will be aimed at modelling the effect of temperature on clock drift.

Sensing One of the main promises of the Future Internet is to bridge the gap between the computing and the real world. To fulfil this vision, a network must have the capability to sense reliably. Temperature can affect the operation of key sensing components such as analog-to-digital convertors (by introducing noise to the measurements). Similarly to the above two pillars, the aim of RELYonIT is to model this negative phenomena in a simple and generic way.

Energy consumption In several of the scenarios envisioned for sensor network systems, the network will be left unattended and without a constant power supply. Due to this reason, energy efficiency has been studied extensively. The extra energy expenditure caused by packet loss (due to interference) is relatively well understood, but the effect of temperature on the overall lifetime of the network has not been studied in the same thorough manner. Our efforts are aimed at this latter direction.

1.1 Our progress thus far

The bird's eye view of our progress is captured in the tables below. The darker the color of the cell, the greater the progress. During the first eight months, our research work has focused more on platform models than on environmental models, and more on temperature than on interference. The reasons for these decisions are twofold.

Focus on temperature The research community has paid significantly less attention to the effect of temperature than to the effect of interference. We therefore saw more potential for novel contributions by focusing on temperature first.

Focus on platforms During the first months we focused our attention on developing a temperature testbed. With this testbed at hand, we thought that it was more important to have first a *general* understanding about the effect of temperature on different platforms, and then delve into more *detailed* effects (environmental models). We also believe that the protocol models will determine the characteristics of the environment that we should focus on (and protocol modelling is a subsequent task).

Table 1.1: Environmental Models

Temperature	Interference

Table 1.2: Platform Models

	Impact of Temperature	Impact of Interference
Signal strength		
Time synchronization		
Sensing accuracy		
Energy consumption		

Overall, the concrete contributions of our work on Tasks 1.1 and Task 1.2 are:

- The design and implementation of a temperature-controlled testbed specifically designed to test sensornet hardware platforms.

- A thorough evaluation of the impact of temperature on the signal strength of multiple hardware platforms.
- A generic platform model to capture the impact of temperature on the received signal strength and on the transmission range of sensornet hardware.
- A thorough evaluation of the impact of temperature on clock drift.
- Initial insights on the effect of temperature on sensing accuracy and energy consumption.
- An initial model for environmental temperature.
- Initial platform and environmental models for interference effects.

2 Temperature - Platform Models

Environmental temperature is known to largely affect the communication performance of wireless sensor and actuator networks (WSAN), but existing work often neglects its impact. For example, temperature can affect clock drift, crystal oscillator startup time, battery capacity and discharge, as well as the performance of the radio transceiver. In a WSAN exposed to harsh environmental conditions, daily or hourly changes in temperature can dramatically reduce the throughput, increase the delay, or even lead to network partitions.

Since temperature has a specific impact on all hardware components used to build sensor nodes, different hardware platforms will be affected in a different way. Existing datasheets describe to some extent the impact of temperature on the component performance, but the existing data is often not detailed enough to derive platform models. For example, the datasheet of the CC2420 transceiver (a common component used in current sensor nodes [41]) only states the temperature range in which the transceiver is operational but it does not state how link quality and transmission error rates vary as a function of temperature within the allowed temperature range.

In this section, using experimental facilities from WP4, we develop platform models that capture how temperature affects the signal strength of low-power communications (Section 2.1) and consequently the topology of a network (Section 2.2), clock drift (Section 2.3), sensing accuracy (Section 2.4), and energy consumption (Section 2.5).

2.1 Impact of temperature on signal strength

Temperature is known to have a significant effect on the performance of radio transceivers: the higher the temperature, the lower the quality of links. Many studies describing experiences from WSN outdoor deployments have reported that diurnal (day/night) and seasonal (summer/winter) fluctuations of ambient temperature have a strong impact on communication quality. Lin et al. [28] have found a daily variation in the received signal strength (RSS) of up to 6 dBm, with the highest RSS values being recorded during night-time. Similarly, in their deployment in an Australian outdoor park, Sun and Cardell-Oliver [39] have measured on-board temperature daily variations between 10 and 50 °C, and noticed that links perform very differently between day and night. Also Thelen et al. [47] have noticed a drastic decrease of RSS at high temperatures in their potato-field deployment.

Results by Bannister et al. [6] from an outdoor deployment and from experiments in controlled scenarios have revealed that an increase in temperature causes a specific reduction in RSS. In their experiments in a climate chamber, the authors observe a linear decrease in RSS of about 8 dB over the temperature range 25-65 °C and show that this reduction may have severe consequences on the connectivity of a network. These results were confirmed by experiments by Boano et al. [9], [8], showing that one can safely decrease the transmission power of communications at low temperatures without deteriorating the performance of the network.

A recent long-term outdoor deployment by Wennerström et al. [48] has further shown that the average packet reception rate (PRR) in a WSN of 16 Tmote Sky nodes dropped by more than 30% when changing temperature from -5 to 25 °C, and that a clear degradation in PRR and average link quality occurred during summer, confirming that daily and seasonal fluctuations of ambient temperature have a strong impact on the quality of sensor network communications.

These existing works simply report the degradation of signal strength and link quality as a consequence of an increase in ambient temperature and do not provide a deeper analysis of the problem. In addition, every reported analysis is unique in terms of experimental setup and hardware. The used radio chips range from Nordic NRF903 [39] and CC1000 [47] to the popular CC1020 [8] and CC2420 transceivers [28], [48], making it difficult to separate general from hardware-specific effects. Bannister et al. [6] have attempted to quantify the loss of RSS due to temperature changes, but only for a limited temperature range and for a single radio chip. Furthermore, when simulating the reduction of communication range and connectivity degradation due to an increase in ambient temperature, the authors assume that communicating nodes have similar temperatures.

Therefore, we go beyond existing work and study the impact of sender and receiver temperature on link quality systematically using different hardware platforms. First, we study the evolution of link quality over one year in an outdoor deployment in Sweden. Our analysis shows that temperature has a strong impact on communication, with visible daily and seasonal differences. Building on top of these results, we carry out a large set of experiments in controlled settings, where we can repeat and alter the conditions at different nodes separately. In all our experiments, we analyse the impact of temperature by measuring the hardware-based link quality metrics in IEEE 802.15.4 compliant radio transceivers [3], namely the received signal strength indicator upon packet reception (RSSI) and in absence of packet transmissions (noise floor), and the link quality indicator (LQI)¹.

2.1.1 Long-Term Outdoor Deployment

We study the evolution of link quality in a sensor network comprising 16 TelosB sensor nodes deployed outside Uppsala, Sweden, in an open field isolated from human activity and absence of electromagnetic interference. Sensor nodes are mounted on poles along a 80 meter straight line at intervals of 0, 20, 40, and 80 meters: on each pole, two nodes are mounted at 0.5 and 1.5 meters height, respectively. The nodes are powered via USB and attached to a Sensei-UU testbed [31], ensuring reliable and continuous data logging.

The software running on the sensor nodes periodically sends packets between every possible pair of nodes and works as follows. Each node is assigned the sender-role in a round-robin fashion every 30 seconds. During this phase, the designated sender transmits one packet per second addressed to each of the other nodes, again in a round-robin manner. When a packet is received by the intended recipient, a response packet addressed to the sender is sent. Each time a sensor node receives a packet – including when it is not the intended recipient – it logs several statistics about the received packet, namely RSSI, LQI, and noise floor. Please notice that the

¹Part of the material used in this section has been submitted for publication [12].

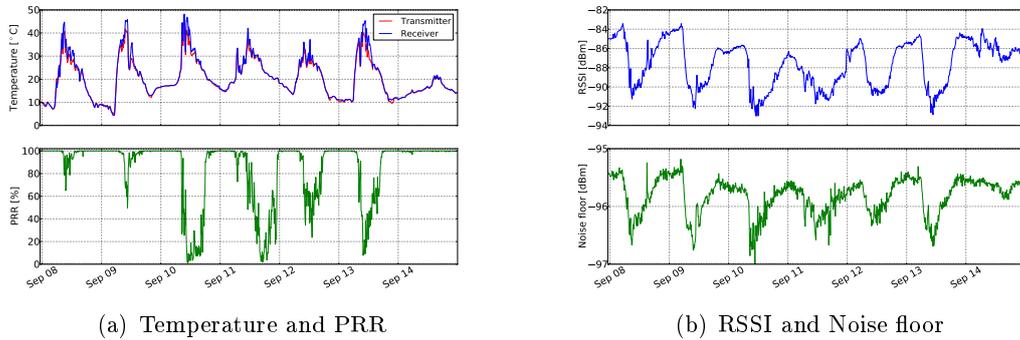


Figure 2.1: Temperature has a strong impact on the quality of links in our outdoor WSN. During daytime, when temperature is high, there is a significant reduction in PRR (a). Also the trend of RSSI and noise floor resembles the one of temperature, with a sharp decrease when temperature increases (b). Values are averaged over a timespan of 10 minutes.

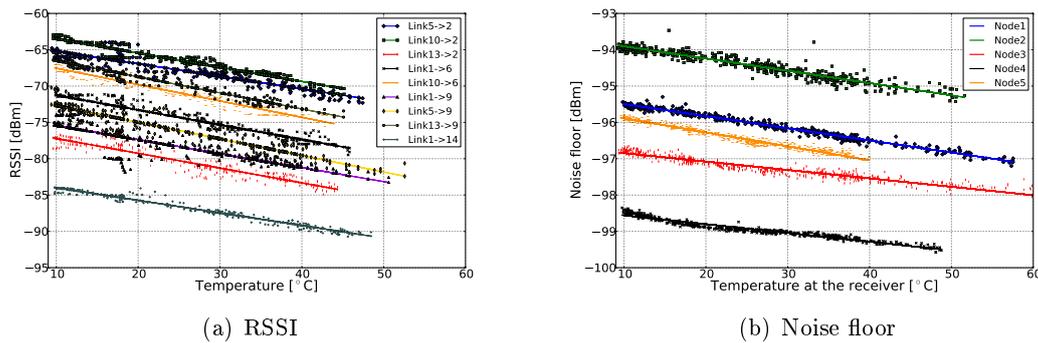


Figure 2.2: The relationship between RSSI and temperature (a) and between noise floor and temperature (b) can be approximated as a linear function, and the trend is similar for different nodes.

RSSI readings from all sensor nodes employed in our experiments are uncalibrated. On-board ambient temperature is measured on each node every two seconds using the on-board SHT11 temperature sensor. More details on the experimental setup can be found in [48].

Impact of temperature on PRR. To highlight the impact that ambient temperature has on the links deployed in our outdoor WSN, we focus on a specific link, close to the edge of the communication range. Fig. 2.1(a) (top) shows the temperature of two nodes (transmitter and receiver) forming a unidirectional link during a week in September. Temperature varies as much as 40 °C between day and night since sensor nodes are enclosed into air-tight enclosures

	1 year	1 month	1 day	1 hour
Lowest temp. (°C)	-22.2	-3.0	7.2	21.2
Highest temp. (°C)	61.3	63.7	63.8	55.9
Temp. difference	82.5	66.7	56.6	34.9

Table 2.1: Largest temperature variations on a single node as seen in our outdoor deployment.

and exposed to direct sunlight. Therefore daily temperature fluctuations may cause a combined overall variation between the two nodes of up to 80 °C. Although the highest variations occur over the 24-hours, temperature can fluctuate by as much as 34.9 °C within one hour, as shown in Table 2.1, in which the largest temperature ranges observed during the 12-months deployment for different time intervals are summarized.

Fig. 2.1(a) (bottom) further shows that each substantial increase in temperature (typically occurring during daytime) results in a decrease in PRR, leading to an almost complete disruption of the connectivity between the two nodes.

Impact of temperature on RSSI and noise floor. The decrease in PRR is strongly correlated with a decrease in the RSSI computed over the received packets, as shown in Fig. 2.1(b) (top), hinting that the change in temperature – and not external interference – was the cause of the packet loss. In particular, the RSSI fluctuates between -84 and -92 dBm, the latter being the threshold below which no packets are received. Interestingly, also the noise floor follows a trend similar to the RSSI and decreases as temperature increases, but to a much lower extent, as shown in Fig. 2.1(b) (bottom).

The strong correlation between temperature, RSSI, and noise floor is highlighted in Fig. 2.2(a) and 2.2(b), respectively. Fig. 2.2(a) shows the RSSI and the combined temperature of sender and receiver for nine links with different link quality over a timespan of three days. The relationship between temperature and RSSI can be approximated as a linear function and is clearly visible despite the intrinsic noise produced by long-term measurements. Using linear regression we have observed that different links have a similar trend, with an average slope of -0.205 and a standard deviation of 0.026.

Fig. 2.2(b) shows the noise floor of five nodes over the same 3 days. Also in this case, the relationship with temperature is approximately linear, with a similar slope among different nodes, but with a less pronounced decrease compared to RSSI (average slope of -0.034 ± 0.006).

2.1.2 Controlled Testbed Experiments

To get a deeper understanding of the effects observed in Sect. 2.1.1, we use experimental facilities from WP4, and augment existing sensor network testbed with the ability of varying the on-board temperature of sensor motes and *reproduce the impact of temperature on link quality in a repeatable fashion* (for further details, please refer to Deliverable 4.2). We use this low-cost testbed infrastructure to systematically study the impact of temperature on different hardware platforms and to isolate the effects of temperature on transmitting and receiving nodes.

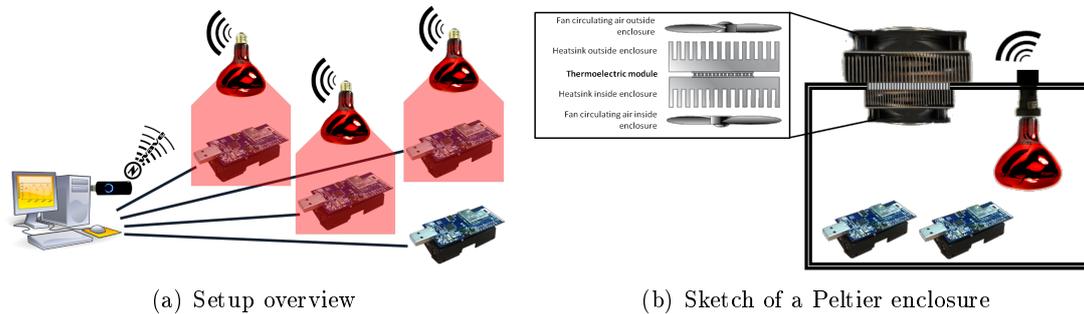
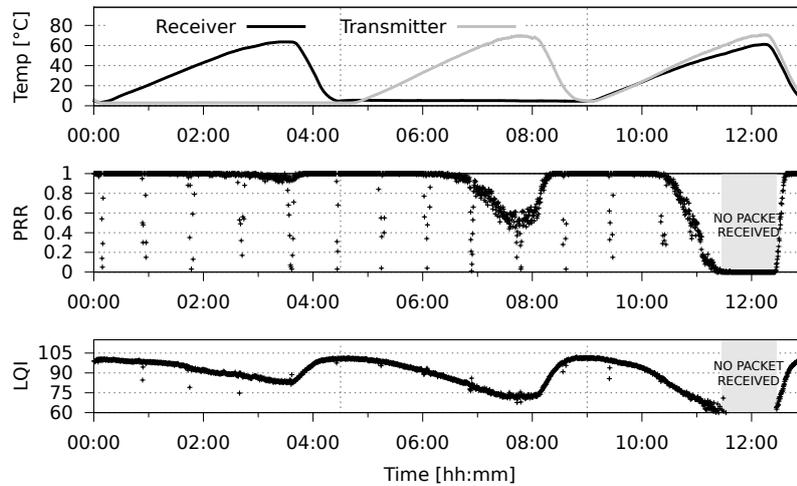


Figure 2.3: Experimental setup in controlled testbed experiments.

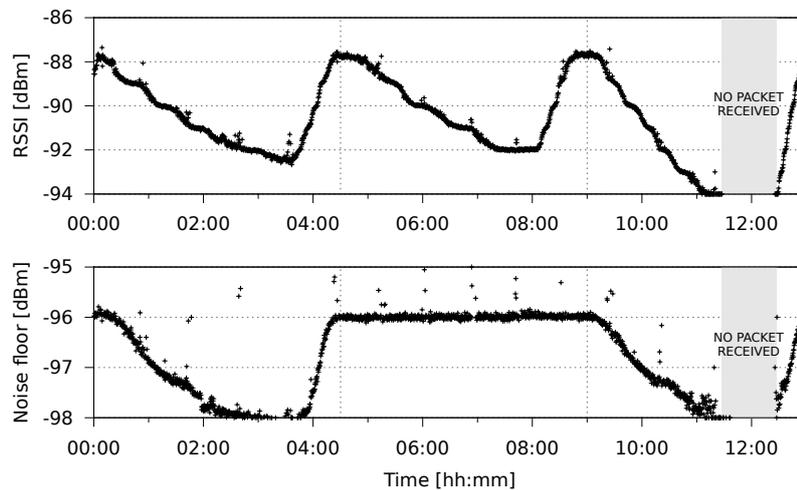
Fig. 2.3(a) shows an overview of our controlled experimental setup. We have extended an existing WSN testbed with the ability of varying the on-board temperature of sensor motes in the range -5 to $+80$ °C using infrared light bulbs placed on top of each sensor node. The light bulbs can be remotely dimmed using the 868 MHz frequency, and hence their operations do not interfere with the communications between the wireless sensor nodes, as the latter use the 2.4 GHz ISM band. In order to cool down the motes below room temperature, we have built custom Polystyrene enclosures as shown in Fig. 2.3(b), in which, in addition to the light bulb, a Peltier air-to-air assembly module by Custom Thermoelectric cools the temperature down to -5 °C when the enclosure is kept at room temperature and the light bulb is off. As we only have a limited number of Peltier enclosures, some of the nodes in the testbed are only warmed by the infrared light bulbs between room temperature and their maximum operating temperature range.

Our testbed is composed of Maxfor MTM-CM5000MSP and Zolertia Z1 nodes employing the CC2420 radio [41], as well as of Arago Systems WisMotes employing the CC2520 transceiver [42]. Sensor nodes are divided in pairs and form bidirectional links operating on different physical channels to avoid internal interference. All sensor nodes run the same Contiki software: each sensor node continuously measures the ambient temperature and relative humidity using the on-board SHT11 or SHT71 digital sensors, and periodically sends packets to its intended receiver at a speed of 128 packets per second using different transmission power levels. Statistics about the received packets are logged using the USB backchannel and are available remotely.

Validation of our controlled setup. Using our controlled testbed setup, we are able to reproduce the impact of temperature on link quality in a very fine-grained way. In a first experiment using Maxfor nodes, every link in the testbed is exposed to three heat cycles. First, each individual node, i.e., first the transmitter and then the receiver, is heated from 0 up to 65 °C. Afterwards, both nodes are heated in the same temperature range at the same time. Fig. 2.4(a) illustrates the impact of temperature on PRR and LQI on a particular link. The evolution of temperature at the transmitter and at the receiver over the 13-hours experiment is shown in the top figure. In correspondence to each increase of temperature, PRR and LQI decrease significantly, with the highest impact occurring when both nodes are heated. With



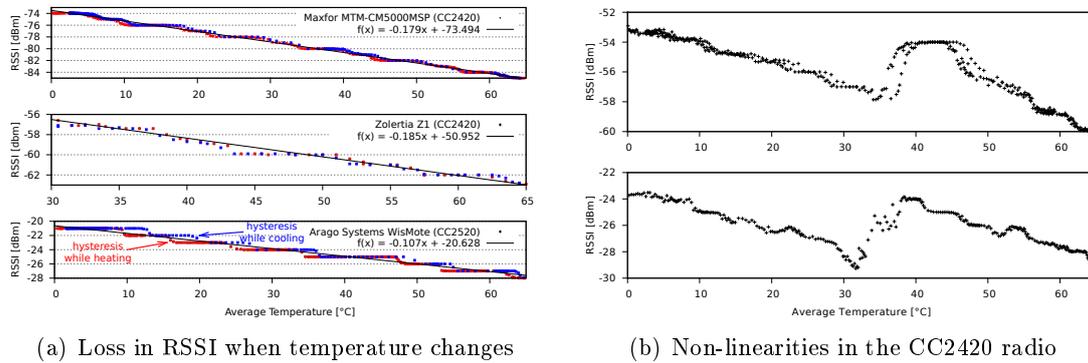
(a) PRR and LQI



(b) RSSI and Noise floor

Figure 2.4: Impact of temperature on the quality of links in our controlled testbed. We heat transmitter and receiver nodes separately first, and then both of them at the same time. When temperature increases, PRR, LQI, and RSSI decrease significantly, with the highest impact occurring when both nodes are heated at the same time. The periodic noise is due to a Wi-Fi access point beaconing in proximity of the testbed.

both nodes heated, indeed, no packet was received and the connectivity between the two nodes was interrupted until the temperature started to decrease. Fig. 2.4(a) also shows that the packet loss rate is more pronounced when the transmitter is heated compared to the case in which only the receiver is heated, something that we have observed in the majority of links in



(a) Loss in RSSI when temperature changes

(b) Non-linearities in the CC2420 radio

Figure 2.5: Figure (a) shows that the relationship between RSSI and temperature is similar when using different hardware platform and can be approximated as a linear function, but with different parameters. Figure (b) shows the non-linearities in the response of the CC2420 radio measured using Maxfor nodes. Temperature on the x-axis is computed as the average temperature of the transmitter and receiver temperature.

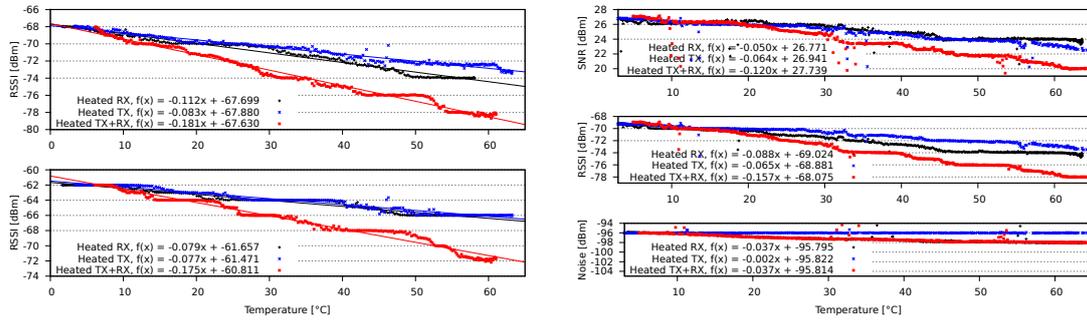
our testbed.

Fig. 2.4(b) illustrates the impact of temperature on RSSI (top figure) and noise floor (bottom figure). The RSSI decreases in a similar way when transmitter and receiver are heated separately, whereas the decrease is more pronounced if both transmitter and receiver are heated at the same time. This proves that temperature decreases both the transmitted and received power [6], whereas the noise floor only decreases when the receiver node is heated, with an absolute variation smaller than the one of RSSI.

These results hence prove the validity of our setup and confirm the measurements obtained in our outdoor deployment, quantifying precisely the impact on temperature on each individual node. We now derive a set of observations obtained running experiments using the same experimental setup, i.e., three heat cycles in which each node is heated individually first and then both nodes are heated at the same time, on different hardware platforms.

The decrease in RSSI is consistent among different platforms. The trend observed in our outdoor deployment showing that RSSI decreases in an approximately linear fashion with temperature holds for different platforms and different radio chips, but with a different slope. Fig. 2.5(a) shows the relationship between RSSI and temperature obtained on different platforms when heating both nodes at the same time. The hardware platforms employing the same CC2420 radio exhibit approximately the same slope.

The decrease in RSSI does not depend on how quickly temperature changes. In our setup, the heat cycles are characterized by a slow increase in temperature followed by a quicker cooling phase, as can be seen in Fig. 2.4(a). This allows us to observe that both RSSI and noise floor are not affected by how quickly temperature varies. Hence, the impact of temperature can be modelled using the absolute temperature value at the transmitter and receiver



(a) Loss in RSSI when using different TX powers (b) Loss in noise floor, RSSI, and SNR for a given link

Figure 2.6: Relationship between RSSI, noise floor, SNR and temperature when transmitter (blue) and receiver (black) nodes are heated individually, and when both nodes (red) are heated at the same time.

nodes.

Discrete steps. On close inspection in Fig. 2.5(a), one can observe discrete steps in the relationship between RSSI and temperature. For the CC2420 platforms, the size of the prominent steps is 2 dBm, whereas for platforms employing the CC2520 radio the step is 1 dBm large. Bannister [5] has attributed the loss of RSSI to the loss of gain in the CC2420 Low Noise Amplifier (LNA). Our experiments bring further evidence to strengthen this claim, as there are references to 2 dBm steps in the CC2420 datasheet [41] with regard to the operation of the Automatic Gain Controller (AGC).

Hysteresis. Fig. 2.5(a) also shows an hysteresis in the relationship between RSSI and temperature that can be seen comparing the RSSI curve obtained when heating and when cooling down the motes. As for the discrete steps, the hysteresis also can be attributed to the operation of the AGC in the CC2420 radio. According to the CC2420 datasheet, hysteresis on the switching between different RF front-end gain modes is set to 2 dBm [41].

Non-linearity in the CC2420 curve. In our experiments, we have also noticed visible non-linearities when the RSSI is ≈ -28 and -58 dBm in the CC2420 platform, as shown in Fig. 2.5(b). These non-linearities were also measured by Chen and Terzis [13], and may lead to a false approximation in case the RSSI of the considered link falls *exactly* in this region (as in the experiments of [6]). When deriving our linear approximation for the CC2420 transceiver, we hence do not consider links falling in this range.

RSSI loss on transmitter and receiver. Fig. 2.6(a) shows the relationship between RSSI and temperature obtained on Maxfor nodes when transmitter and receiver nodes are heated individually and when both nodes are heated at the same time. Top and bottom figures refer to the same link, but are obtained using a different transmission power. Despite the link is

the same, the relationship between RSSI and temperature is slightly different, with a steeper decrease when the receiver is heated in the top figure. Although a comparison between curves is difficult due to the AGC operations (depending on whether we capture the transition between two discrete steps, we may obtain slightly different slopes), by averaging the data from all our experiments we have obtained a relationship between receiver and transmitter of 0.5348 ± 0.061 . The RSSI seems hence to have a slightly steeper slope when the receiver node is heated.

Impact on noise floor and SNR. Fig. 2.6(b) illustrates how noise floor, RSSI, and signal to noise ratio (SNR) vary on a given link when transmitter and receiver nodes are heated individually and at the same time. Since the noise floor decreases only when the receiver is heated, an increase in temperature on the transmitter has a higher impact on the SNR compared to an increase in temperature at the receiver. This also explains the different impact in PRR when heating the nodes individually that we observed in Fig. 2.4(a).

2.1.3 Generalized model of the effect of temperature on RSS

The effect of temperature on electric conductors and semiconductors is well known. Various models have been created for a large range of devices to capture the relation between ambient temperature and electric conductance (and current leakage). Our goal is to build on top of this knowledge to create a generic model for low-power radio transceivers. It is important to remark that the goal of our model is not to benchmark a specific radio chip against others, as this is already done by manufacturers. Our goal is to develop a simple model to predict the performance of a network under extreme environmental settings. We now describe the overarching effect of temperature on radio transceivers and derive a generic model for low-power wireless transceivers.

In electric conductors, a higher temperature increases the resistance of the medium, whereas in semiconductors it leads to current leakages. In practice this means that, for a given voltage, a higher temperature reduces the current and hence the power of a device. In radio transceivers, these phenomena imply that a raise in temperature will reduce the SNR. A decrease in SNR leads to a lower link quality and a shorter radio link, which in turn may lead to lower throughput, higher delay or even network partitioning. Hence, our goal is to model the effect of temperature on SNR. Denoting PL as the path loss between a transmitter-receiver pair, P_t as the transmission power, P_r as the received power, and P_n as the noise floor at the receiver, the SNR is known to be:

$$\begin{aligned} SNR(dB) &= P_t - PL - P_n \\ &= (P_t - P_n) - (P_t - P_r) \end{aligned} \quad (2.1)$$

As we have shown in our empirical measurements, an increasing temperature has 3 main effects on the signal strength of radio transmissions; it (i) decreases the transmitted power, (ii) decreases the received power, and (iii) decreases the noise floor. We now model these three effects in Eq. 2.1.

A first-order model Denoting α , β , γ as constants with units dB/K , and T_t , T_r as the temperature in Kelvin of transmitter and receiver, the effect of temperature on SNR can be

defined as:

$$\begin{aligned}
 SNR &= (P_t - \alpha\Delta T_t) - (PL + \beta\Delta T_r) \\
 &\quad - (P_n - \gamma\Delta T_r + 10\log_{10}(1 + \frac{\Delta T_r}{T_r})) \\
 &= P_t - PL - P_n - \alpha\Delta T_t \\
 &\quad - (\beta - \gamma)\Delta T_r - 10\log_{10}(1 + \frac{\Delta T_r}{T_r})
 \end{aligned} \tag{2.2}$$

The proportional relation between ΔT and the constants α (effect on transmitted power), β (effect on received power) and γ (effect on noise floor) is based on the empirical observations made in the previous sections. The term $10\log_{10}(1 + \frac{\Delta T_r}{T_r})$ is derived analytically from the well-known thermal equation. There are two important trends to highlight in this model. First, changes in temperature have a higher impact on the transmitted and received powers (linear relation of α and β), than on the thermal noise (logarithmic relation). Second, to some extent it is counter-intuitive that a higher temperature decreases the noise floor (negative sign of γ). This effect was also observed by Bannister, and he hypothesizes that it is due to the losses in the signal amplifier [5]. That is, a higher temperature not only reduces the gain of the signal but also the gain of the noise, and hence, the received signal strength (RSSI) is lower for both.

The accuracy of our model depends on identifying the right values for α , β and γ . In our case, these parameters are given by the slopes of the linear trends observed in our empirical results. These parameters are platform dependant, and hence require a systematic and fine-grained evaluation. Our testbed was designed to accomplish exactly that. For example, a network manager willing to deploy a network using the Maxfor platform, can use the slopes obtained in Fig. 2.6(b): $\alpha = 0.065$, $\beta = 0.088$ and $\gamma = 0.037$. Assuming that the network will be deployed in an environment where the maximum and minimum day temperature are 50 and 5°C respectively, the network manager can predict that the links can suffer an attenuation of $(\alpha + \beta - \gamma)\Delta T = 5.22$ dB (5 dB according to the SNR measurements in Figure 2.6(b) top). This level of attenuation can easily push a good link (with 100% PRR) to have a PRR of 0%.

2.2 Impact of Temperature on Radio Range Shrinkage and Network Topology

Our experiments have clearly established that the Received Signal Strength (RSS) decreases with an increase of the temperature. The numerical value of this reduction itself does not highlight its impact on the communication. In this section we investigate the impact of the temperature on the network connectivity in terms of the radio range and the network topology.

2.2.1 Radio Range Shrinkage and Loss of Links Due to Temperature

The received signal strength (RSS) decreases with the increase of temperature. A lower RSS has a direct impact on the transmission range of nodes. We analyzed the impact of temperature on the transmission range.

We base our analysis on the *Log-Distance Path Loss Model*. According to this model, the path loss at distance d is

$$P_t - P_r = PL_0 + 10\eta\log_{10}\frac{d}{d_0} + X_g$$

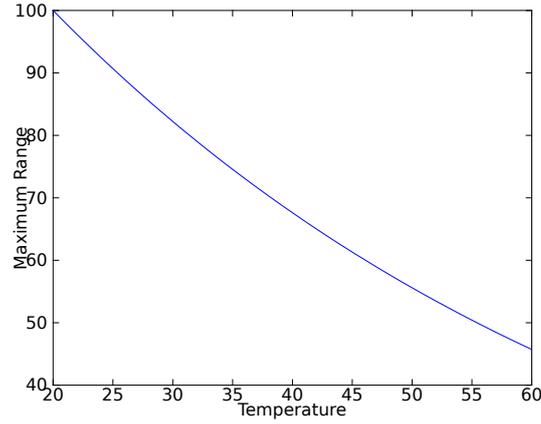


Figure 2.7: The range shrinks with the increase of temperature.

P_t is the transmitted power in decibel, and P_r is the received power in decibel. PL_0 is the path loss (in decibel) at a reference distance d_0 , η is the path loss exponent, and X_g is the attenuation due to fading. In this analysis we assume that $X_g = 0$.

$$P_r = P_t - PL_0 - 10\eta \log_{10} \frac{d}{d_0}$$

Assume that the minimum power required to detect the signal at the receiver (receiver sensitivity) is P_{min} , and this power is at d_{max} distance from the receiver.

$$P_{min} = P_t - PL_0 - 10\eta \log_{10} \frac{d_{max}}{d_0} \quad (2.3)$$

We know from our experiments that the power at the receiver goes down further because of the increase in temperature. This loss is proportional to the *increase* (Δt) in the temperature. Assume that all the nodes are heated uniformly for the following discussion. We can modify the above equation to take this into account.

$$P_{\Delta t} = P_t - PL_0 - 10\eta \log_{10} \frac{d_{max}}{d_0} - k\Delta t$$

$P_{\Delta t}$ is the new signal power at the receiver, and k is the gradient of the best fit line for our traces.

It is obvious that $P_{\Delta t} < P_{min}$. Therefore, d must be reduced to $d_{\Delta t}$ in order to achieve P_{min} .

$$P_{min} = P_t - PL_0 - 10\eta \log_{10} \frac{d_{\Delta t}}{d_0} - k\Delta t \quad (2.4)$$

From 2.3 and 2.4,

$$\begin{aligned} P_t - PL_0 - 10\eta \log_{10} \frac{d_{max}}{d_0} &= P_t - PL_0 - 10\eta \log_{10} \frac{d_{\Delta t}}{d_0} - k\Delta t \\ -10\eta \log_{10} \frac{d_{max}}{d_0} &= -10\eta \log_{10} \frac{d_{\Delta t}}{d_0} - k\Delta t \\ k\Delta t &= 10\eta \log_{10} \frac{d_{max}}{d_0} - 10\eta \log_{10} \frac{d_{\Delta t}}{d_0} \\ k\Delta t &= 10\eta \log_{10} \frac{d_{max}}{d_0} \frac{d_0}{d_{\Delta t}} \\ k\Delta t &= 10\eta \log_{10} \frac{d_{max}}{d_{\Delta t}} \\ \frac{d_{max}}{d_{\Delta t}} &= 10^{\frac{k\Delta t}{10\eta}} \\ d_{\Delta t} &= d_{max} * 10^{-\frac{k\Delta t}{10\eta}} \end{aligned} \quad (2.5)$$

Figure 2.7 depicts the loss of range caused by the increase in temperature. The initial range is set to 100 units at the temperature of 20° C. $k = 0.17$, which is the average tangent for the RSSI vs temperature plots generated from one of the controlled experiments. The temperature range (20° – 60° C) and the path loss exponent ($\eta = 2$) were selected to match the similar plot in Bannister's thesis [5]. However, note that only the difference in the temperature matters.

This is quite similar to the Figure 12 in Bannister's thesis. However, Bannister mentions that the results were obtained using a simulation. In our approach these results were obtained using an analytical model. In addition, Bannister has used values for the receiver sensitivity, output power, and the path loss at the reference distance as parameters to his model. The Equation 2.5 clearly indicates that the range is independent of those parameters. The range after an increase of the temperature by ΔT , only depends on the range before the increase (d_{max}), the parameter k , and the path loss exponent (η). This is an important distinction since the only platform-dependent parameter required for this analysis is k , which can be found through the experiments mentioned before.

The reduction in the effective radio range is quite significant: it is reduced by almost 55% for an increase of 40°C.

To analyse the effect of the temperature on the number of links in a network topology, we assume that the nodes are distributed evenly in a two-dimensional plane with a node density of ρ nodes per unit square. Therefore the number of nodes within the d_{max} range is $\rho\pi d_{max}^2$. We can take this to be the number of direct links that the node in consideration has to other nodes.

When the temperature is increased by Δt , the number of other nodes within the range is reduced to $\rho\pi d_{\Delta t}^2$. Hence, the fraction of links lost because of an increase in Δt° C, $L(\Delta t)$ is proportional to the reduction in the area covered by the radio range.

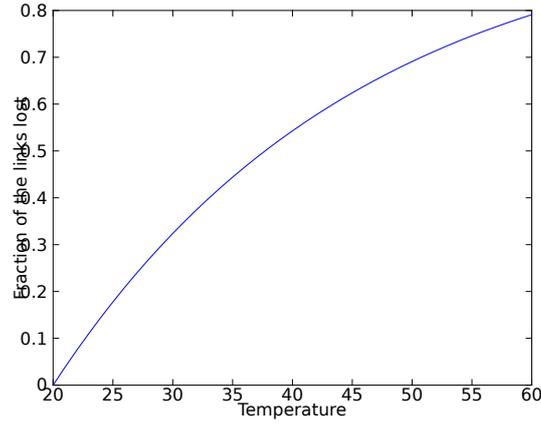


Figure 2.8: A node loses links with the increase of temperature.

$$\begin{aligned}
 L(\Delta t) &= \frac{\rho\pi d_{max}^2 - \rho\pi d_{\Delta t}^2}{\rho\pi d_{max}^2} \\
 &= 1 - \frac{\Delta t^2}{d_{max}^2}
 \end{aligned}$$

We know from the Equation 2.5 that $\frac{\Delta t}{d_{max}} = 10^{-\frac{k\Delta t}{10\eta}}$.
 Therefore,

$$\begin{aligned}
 L(\Delta t) &= 1 - \left(10^{-\frac{k\Delta t}{10\eta}}\right)^2 \\
 &= 1 - 10^{-\frac{k\Delta t}{5\eta}}
 \end{aligned} \tag{2.6}$$

Figure 2.8 depicts the fraction of the links lost with the increase of the temperature. Again, $k = 0.17$ is the average gradient as mentioned before, and η is taken as 2.

Almost 80% of the links are lost for an increase of 40° C. This result is similar to that obtained by Bannister (Figure 14 in the thesis [5]) regarding node connectivity.

We assumed, for the sake of simplicity, that both the sender and the receiver have the same temperature in deriving the equations 2.5 and 2.6. However, by using the model presented in the previous section and assuming that $10\log_{10}(1 + \Delta T_r/T_r) = 0$, we can relax this restriction. This results in the following more generic models for the link loss and the reduction of the range.

$$d_{\Delta t} = d_{max} * 10^{-\frac{\alpha \Delta T_t + (\beta - \gamma) \Delta T_r}{10\eta}} \quad (2.7)$$

$$L(\Delta t) = 1 - 10^{-\frac{\alpha \Delta T_t + (\beta - \gamma) \Delta T_r}{5\eta}} \quad (2.8)$$

This model relies on three platform dependent parameters, α , β , and γ , but it allows us to analyze the effects of changing the temperature at the sender and the receiver independently. For example, we can assume that the temperature is changed only at the sender. Under this assumption the model reduces to the following.

$$d_{\Delta t} = d_{max} * 10^{-\frac{\alpha \Delta T_t}{10\eta}} \quad (2.9)$$

$$L(\Delta t) = 1 - 10^{-\frac{\alpha \Delta T_t}{5\eta}} \quad (2.10)$$

If we use the value for α from the previous section ($\alpha = 0.065$), it is clear that the range is reduced by 26% and the number of links goes down by 45% for an increase of the temperature by 40° C at the sender.

Similarly, if one sender is kept at a constant temperature and the temperature is increased at the receivers (all the other nodes) we obtain the following model.

$$d_{\Delta t} = d_{max} * 10^{-\frac{(\beta - \gamma) \Delta T_r}{10\eta}} \quad (2.11)$$

$$L(\Delta t) = 1 - 10^{-\frac{(\beta - \gamma) \Delta T_r}{5\eta}} \quad (2.12)$$

Again, we can use the β ($= 0.088$), and γ ($= 0.037$) from the previous section as examples. For a 40° C increase at the receivers the range is reduced by 21% and the numbers of links is reduced by 37%.

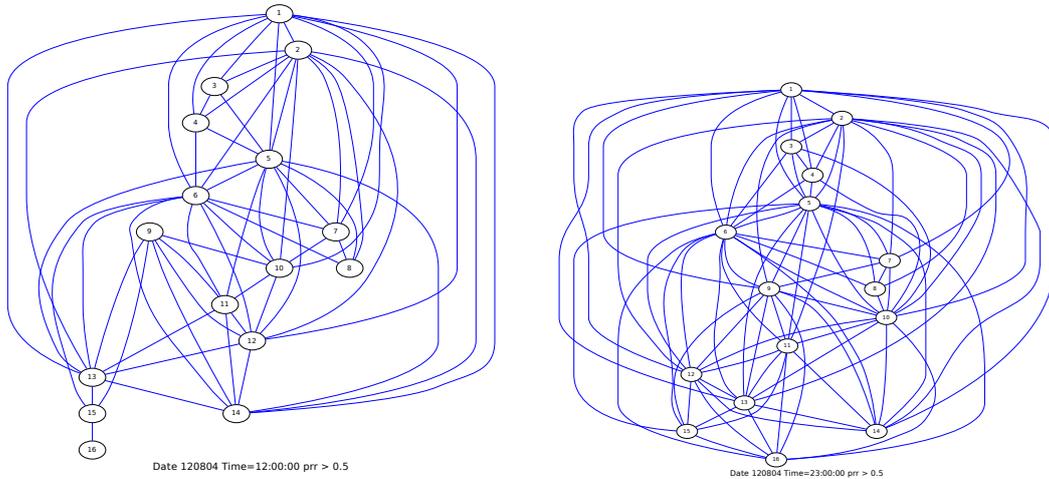
Since α , ($= 0.065$), is quite close to $\beta - \gamma$, ($= 0.051$), the losses related to changing the temperature at only the sender or only the receivers are also quite close. Since there are also errors in estimating these parameters, it would be reasonable to assume (as a convenient approximation) that half the loss is due to the receiver and half is due to the sender for this particular platform used to estimate α , β , and γ .

2.2.2 Impact of Temperature on Network Topology

Temperature changes can considerably affect the network topology and, consequentially, routing protocols. Our path-loss model predicts the loss of range and the loss of links, hence the change in topology, due to the increase of temperature. This theoretical model has shown that the link loss can be considerable even for temperature variations observed within a single day.

We have been able to observe these effects also on network topologies in actual sensor network deployments. We used the data generated from a long-term outdoor sensor deployment with 16 nodes at Uppsala University to analyse how temperature affects the network topology. The experimental setup is the same as the one described in the previous section.

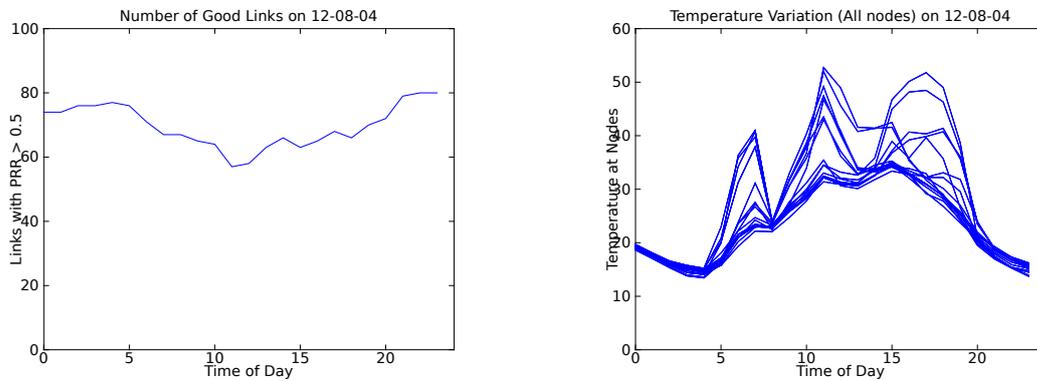
Figure 2.9(a) and 2.9(b) depicts the network topologies at 12:00 hours and 23:00 hours of a particular day. The topology was generated using only the “usable links”; i.e., the links with



(a) Topology at 12:00 hours.

(b) Topology at 23:00 hours.

Figure 2.9: The network topology changes considerably within 11 hours.



(a) The number of usable links.

(b) Temperature variation.

Figure 2.10: The number of usable links goes down with the increase in temperature.

Packet Reception Ratios (PRR) larger than 0.5. We can clearly see that there is a significant change in the network topology during this 11 hour period.

Figure 2.10(a) shows the change of the number of usable links over the 24 hour period and Figure 2.10(b) shows the variation of temperature, as recorded by each of the 16 nodes, over the

same time period. The figures show that the usable links become fewer when the temperature increases during the day. The number of usable links varied between 57 and 80 within this 24 hour period. The network diameter changed from 4 (at 12:00 hours) to 2 (at 23:00) hours. During this time period different nodes have experienced temperatures ranging from 13.4°C to 52.7°C. These changes would have a significant effect on the routing protocols.

2.3 Impact of temperature on clock drift

Of the various components which collectively form an embedded system, some of the most important and fundamental ones are the clocks. All embedded systems from high end super computers, to hand-held devices right down to integrated System-on-Chip (SoC) platforms [17], [45], [46], all use a variety of clocks operating at different frequencies to drive its components for processing, communication, storage etc. Unintended variations in the frequency of these clocks may have unpredictable effects on the operations of system components leading to undesirable consequences with impact upon the applications they support. It is the aim of this section to measure and model how environmental conditions can affect the frequency of system clocks to help better predict and manage the overall effects on the system.

Three types of clocks have been identified in a typical embedded platform [16], [2], [15]. These include:

- The Primary system clock used to control instruction speed in addition to other internal processor features.
- The Real-time clock used to provide system timing and application synchronisation.
- Other Platform clocks such as the radio clock used to drive radio transmission.

Each of these clocks have been subjected to examinations with respect to varying temperature and documented within the following sub-sections.

Primary System Clock. A primary system clock can either be external to a processor or self-contained internally within a processor. Many embedded platforms typically drive their processor clocks from an internal source to reduce the necessary external components of an embedded system. This source will drive the various internal Micro-Controller Unit (MCU) clocks such as the primary clock, which determines the processing speed, as well as other peripheral/auxiliary clocks used for timers, inter-chip communications etc. It is important that the source of these clocks is stable over time, as fluctuations in their frequencies can have wide repercussions and result in possible system failure. Clock instabilities can cause effects such as inconsistent processing speed, communication errors, storage errors to name but a few.

Real-Time Clock. A second platform clock that is common to many system platforms is the Real-Time Clock (RTC). The RTC is typically driven by an external 32.768 KHz watch crystal which gives 2^{15} tick in one second, allowing accurate real-world timing for software operations. The operating system, applications, communication protocols etc. utilise this clock for scheduling their operations. Variances in this clock can lead to undesirable effects on systems,

and can result in malfunctions within communications and application processes.

Other Platform Clocks. Many platforms employ additional on-board clocks to drive components such as storage and communications. These could also be susceptible to temperature effecting the components they drive. It is important to understand how such clocks behave when temperature varies to ensure dependable operation.

One identified common clock that many wireless embedded platforms [16], [2], [15] include is the radio clock used to drive communication. This clock is particularly important as fluctuations in frequency can cause communications failure. For less integrated platforms this clock is often supplied from an external oscillator, however on more integrated platforms such as the ST MB950 this clock is driven from the same source as the primary clock and as such will have the same temperature/frequency profile.

2.3.1 Experimental Results

This sub-section will examine how the three identified clocks sources are effected by temperature.

Primary System Clock

To evaluate how the effects of temperature can impact upon the stability of the primary system clock of an embedded system, a number of embedded platforms are observed whilst varying temperature under controlled conditions. These platforms include the Maxfor MTM-CM5000MSP [29], the Zolertia Z1 [50], and the STMicroelectronics MB950 [38].

Maxfor MTM-CM5000MSP. The first system examined is the Maxfor MTM-CM5000MSP mote, a popular Moteiv Tmote Sky [16] clone. On this platform, the primary MCU is the Texas Instruments MSP430F1611 [43]. Its system and auxiliary clocks are driven from an internal Digital Controlled Oscillator (DCO). The frequency of the DCO is factory calibrated, however, fine grain adjustments to its operational frequency can be made in software by configuring the internal registers, which modifies the DCO frequency by altering its voltage. For the purpose of these experiments, the speed of the clock was fixed to 3.9 MHz. This is the default frequency used by the ContikiOS for this platform.

A simple test application was written for the ContikiOS operating system which toggled a GPIO pin at a fixed rate determined by the frequency of the system clock. Nodes were programmed with the compiled firmware for this application. To control the number of variables for this experiment, all interrupts were disabled. The processor toggled a GPIO pin after a small fixed number of no operation (NOP) operations in a tight loop. These measures ensured that only the frequency of the processor would impact upon the rate at which the GPIO pin was toggled and that no other external factors could obscure these results. The period between the GPIO toggling was measured externally to the platform using a calibrated oscilloscope connected to the GPIO pin. The temperature of the mote was controlled using a temperature lamp, fixed at 10 cm above the node. The power output of the lamp was controlled by a wireless dimmer switch [22] which facilitated 31 power output levels. The power output, and therefore the intensity of the heat from the lamp, was slowly increased over a one hour period

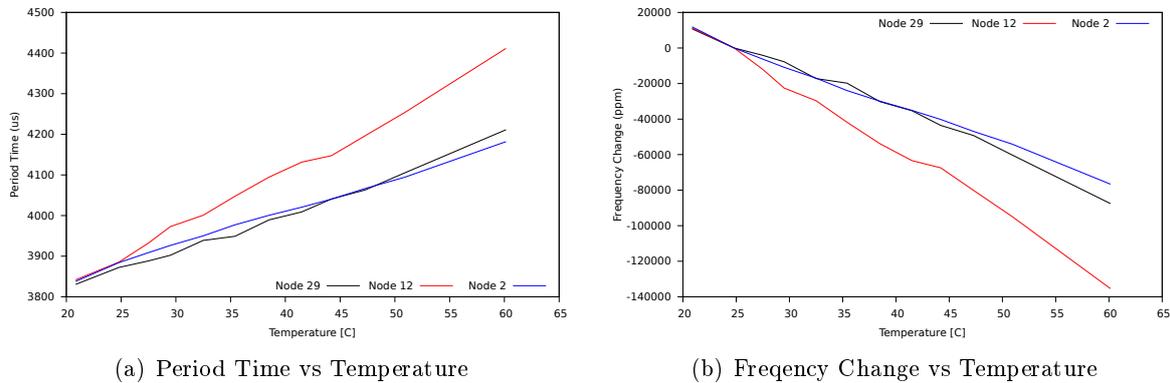


Figure 2.11: Maxfor MTM-CM5000MSP Primary System Clock

by stepping through each of the 31 power setting.

Figure 2.11(a) presents the results of this investigation. The plot contains the results of a test carried out on three identical Maxfor MTM-CM5000MSP motes. The temperature of each node started at ambient room temperature of 21 °C. The GPIO toggle period was found on average to be 3.84 μs for all nodes at this temperature. As the temperature increased, the period also increased accordingly, showing a proportional linear relationship between the period and temperature. At a temperature of 60 °C the period was measured to be 4.21 μs for node 29, 4.41 μs for node 12 and 4.18 μs for node 2. Figure 2.11(b) depicts these results in regards to the frequency change (PPM) demonstrating the change in clock frequency in relation to frequency at room temperature. It is shown that the frequency falls by on average 11.23% after a temperature increase of 40 °C. The rate of change in the period was calculated at -0.257% for each one degree increase in temperature.

Zolertia Z1. The second platform that was examined was the Zolertia Z1 mote [50]. The Z1 can be considered as a modern version of the T-mote offering much of the same or improved capabilities. The Z1 utilises the MSP430F2617 MCU [44], a second generation MSP430 ultra-low power 16-bit MCU from the same family as the MCU used on the T-mote platform. Similarly to the Maxfor mote above, the primary clock of the Z1 is sourced from an internal DCO. To examine the susceptibility to temperature variations of the Z1 DCO, the same test configuration as used previously was applied to this platform to measure the frequency variance over temperature. Only minor changes to the test firmware were applied for platform specific configurations.

Figure 2.12(a) presents the findings for the Z1 platform. The period of a fixed frequency of the primary system clock was again determined by toggling a GPIO pin and its frequency measured externally. The experiment was repeated on two identical Z1 nodes. The results show the same linear relationship between period length and temperature but with a significantly decreased gradient. For node A at ambient temperature of 21 °C, the period was measured at 11.95 μs and at 54 °C was 12 μs . An increase of only 50 ns over a change of 30 degrees Celsius. Less variation was seen with node B where the period only increased from 11.91 μs to 11.93

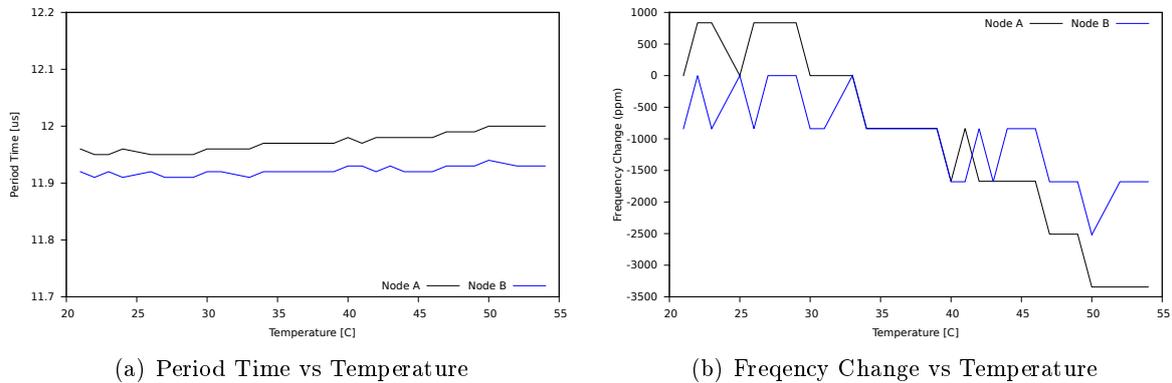


Figure 2.12: Zolertia Z1 Primary System Clock

μs over the 30 degrees range, demonstrating a change of 20 ns. Figure 2.12(a) demonstrates how clock frequency changed over this temperature range. As the temperature increased of the 33 °C range tested, the frequency decreases by an average of 0.25%.

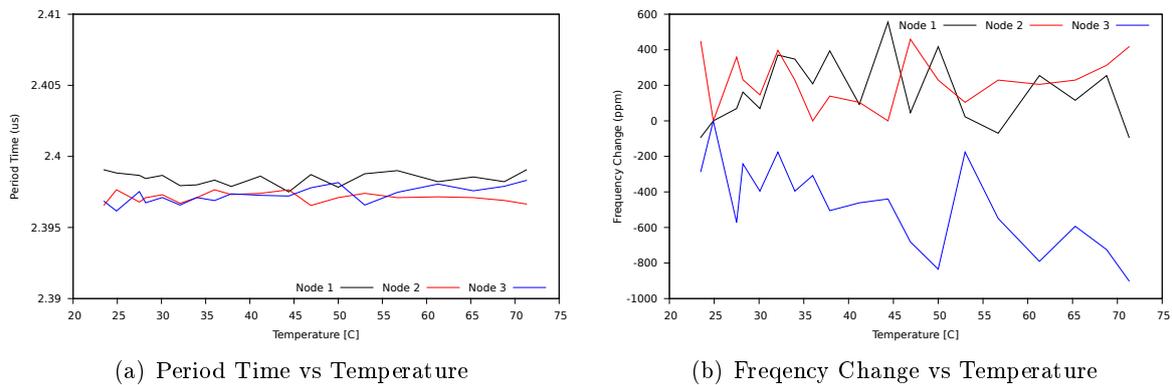


Figure 2.13: ST MB950 Primary System Clock

STMicroelectronics MB950. The final platform investigated was the ST MB950 [38] which utilises an Arm Cortex-M3 STM32W micro-controller [37]. This platform is an example of a modern embedded platform with processor and transceiver integrated and yields a significantly increased processing power over the previously examined platforms. This platform differs to the earlier investigated platforms as it utilises an external 24 MHz oscillator to drive the internal clocks of the MCU. The same experiment was carried out using the ContikiOS firmware and test application, again with only platform-specific changes made to the firmware. The results are shown in Figure 2.13(a) and again appear to indicate a relationship between clock period and temperature, however the clock variance for this platform is minimal. At

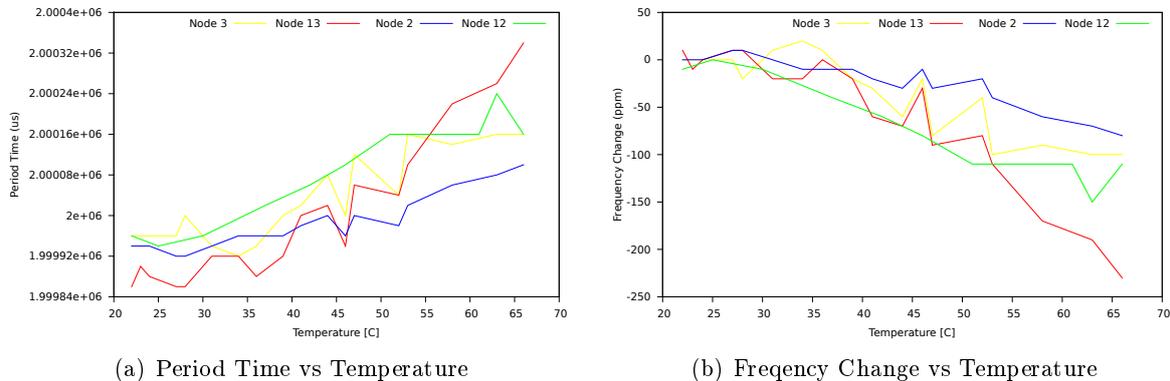


Figure 2.14: Maxfor MTM-CM5000MSP RTC Clock

ambient temperature of 24 °C the average clock period was measured as 2.397 μs and increased slightly to 2.398 μs at 72 °C, demonstrating change of just 1 ns over the tested temperature range. Figure 2.13(a) reflects how clock frequencies changed in regards to the ambient temperature for the temperature range. For these clocks it would appear that the correlation between clock frequency and temperature for this platform is no longer a linear relationship but a form of quadratic relationship. However these clocks would appear to be more resilient against temperature variance with an average frequency increase of just 0.06% increase over the 68 °C tested temperature range.

Real-Time Clock

We examined the Maxfor MTM-CM5000MSP mote to evaluate the effects of temperature changes on the RTC clock. The mote was programmed to raise a GPIO pin for 2^{16} RTC ticks, which is equivalent to two seconds, in a loop whose timing could then be measured by an oscilloscope. The oscilloscope had a sampling resolution of 0.00008 μs . The mote was heated from room temperature to 70 °C similarly to the previous experiment by the use of a heat lamp placed 10 cm above the mote and controlled by a wireless dimmer module, progressively stepping through each of its 31 power output levels.

Figure 2.14(a) presents the result from our experiment. All nodes follow the same trend where the clock was measured faster than the expected frequency at lower temperatures and slower at higher temperature. At 22 °C the lowest measuring point, the nodes were found to be an average of 70 μs fast whilst at the highest temperature of 66 °C the nodes were found to be 190 μs slow. Node 13 had displayed the larger error being 140 μs fast at 22 °C and 340 μs slow at 66 °C. Node 3 reached the target period time of 2 seconds at 39 °C whilst node 13 at 41 °C, node 2 at 44 °C and node 12 at approximately 35 °C. Although there is a relationship between temperature and the period this does not appear to be linear. Figure 2.14(b) presents the same results as a frequency change in ppm from the base frequency measured at 25 °C. Nodes 2, 3 and 12 exhibited a similar trend across all temperatures whilst node 13 had a more extreme gradient. The values ranged from 20 ppm to -460 ppm.

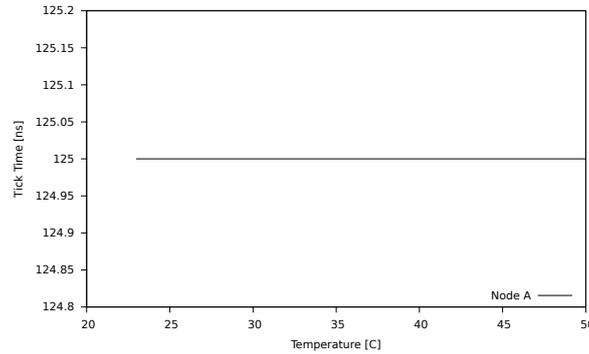


Figure 2.15: CC2420 Radio Clock Period

Other Platform Clocks

On the TI CC2420 transceiver, the radio used on the majority of research platforms including the Tmote/Maxfor mote, the clock is sourced from an external 16 MHz crystal oscillator. To accurately determine the operational frequency of these crystal oscillators, the radio can be placed in a test mode where a divided version (8 MHz) of this clock can be outputted via the CCA pin. Using an oscilloscope connected to this pin, the frequency of the crystal can be measured. The Maxfor mote was configured in this way and was subjected to experiments similar to the previous experiments. A heat lamp was used to vary the node temperature with an oscilloscope connected to the CCA pin to measure how the radio clock is affected.

Figure 2.15 illustrates these findings. It can be determined that the radio clock period is stable at $125ns$ across all tested temperatures. Our conclusion from these results is that the external clock source must compensate for temperature variations and therefore is unaffected as the temperature is varied.

2.3.2 Clock Drift Models

In the previous sub-section two types of clock sources were identified: the RC and external oscillator. Both of these sources were shown to be affected by changes in temperature. The following text will describe a model that approximates those observed effects.

RC-Clock Model. An RC-clock uses resistive and capacitive components whose profiles alter as temperature changes. These profile variations are linear and can be modelled using Equation 2.13. The circuit is designed to have a frequency f_0 close to the target frequency at room temperature T_0 . A_1 is a temperature coefficient with unit $ppm/^\circ C$.

$$f(T)/f_0 = A_1 \cdot (T - T_0) \quad (2.13)$$

We used this model to describe the measured clock changes dependent on the temperature variance for both the Maxfor MTM-CM5000MSP and the Zolertia Z1 (see Section 2.3.1). The results are shown in Figure 2.16(a) and Figure 2.16(b). The determined temperature coefficient for each node is shown in Table 2.2.

It can be seen that the model is unique for each node. It is therefore not possible to determine a generic model which would be applicable for all nodes even of the same type. The model pro-

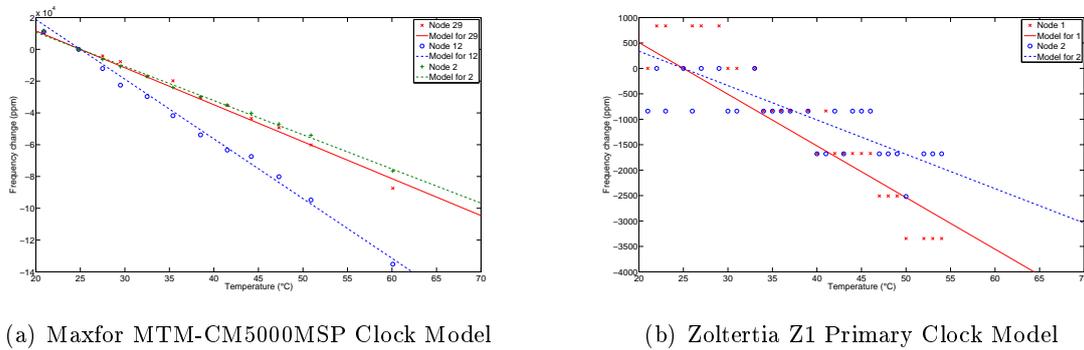


Figure 2.16: Model Fitting for RC Clocks

Node	A_1 [ppm/°C]
Node 12 (Maxfor)	-3756.5
Node 2 (Maxfor)	-2150.6
Node 29 (Maxfor)	-2324.4
Z1 1 (Z1)	-101.56
Z1 2 (Z1)	-67.555

Table 2.2: Determined RC temperature coefficients for the evaluated Maxfor MTM-CM5000MSP and Zoltertia Z1 nodes.

vides a linear approximation to the measured Maxfor and Z1 readings, however jitter is present due to the limited accuracy of the experimental measurements. We consider the residual errors for these experiment which is the measurable error or deviation of the actual data from the statistically obtained estimate of the model. For these experiments there was an average residual of 1919.44PPM for the Maxfor nodes, with a maximum of 5873PPM for node 29 at 61 °C, and an average residual of 462.86PPM with a maximum of 1265PPM for node 1 at 29 °C.

RT-Clock Model. An RT-clock oscillator circuit generally relies on a quartz crystal. The crystal frequency depends on the shape of the quartz (its cut) which changes with temperature. Hence the cut of a quartz has significant impact on the frequency changes dependant on temperature changes.

Crystals are cut to resonate close to the target frequency at room temperature T_0 . Room temperature is generally assumed to be $T_0 = 25^\circ$ by crystal manufacturers. When temperature either increases or decreases the oscillator frequency changes. A number of crystals use a cut (For example, a CT, DT, SL, X, XY, BBT cuts) that can be described by a quadratic function [21]. The frequency of the oscillator changes when temperature increases or decreases. Another commonly used crystal cut is the AT cut which can be described using a cubic function [21]. The different frequency temperature dependencies are schematically shown in Figure 2.17.

The outlined behaviour of different crystal cuts can be modelled using Equation 2.14. f_0

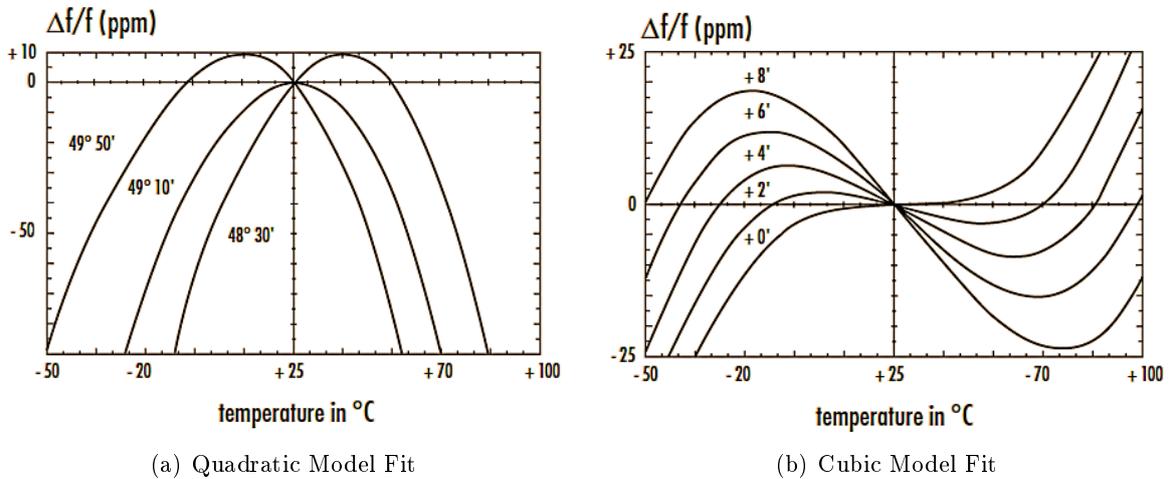


Figure 2.17: Example of two different model that describe Frequency change with Temperature for RT Crystals of different cuts. These figures are found in [21]

is the frequency of the RT-clock at room temperature T_0 . $F(T)$ is the RT-clock frequency at temperature T . A_1 is a temperature coefficient with unit $ppm/^\circ C$. A_2 and A_3 are temperature coefficients with unit $ppm/^\circ C^2$ and $ppm/^\circ C^3$.

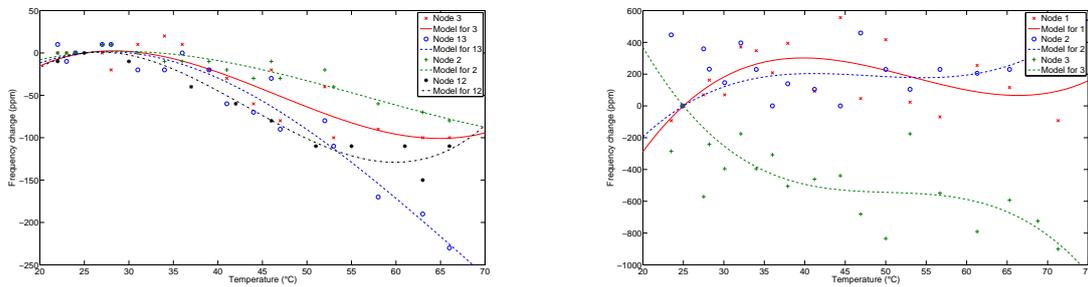
$$f(T)/f_0 = A_1 \cdot (T - T_0) + A_2 \cdot (T - T_0)^2 + A_3 \cdot (T - T_0)^3 \quad (2.14)$$

If it is known that the crystal has a cut that is described by a quadratic function the coefficient A_3 can be set to zero. However, in practice it is not always possible to determine the crystal cut type. Sensor node manufacturers often use different crystals when producing different batches of nodes which may have different cuts. Also, the available datasheets often do not provide this level of detail.

To model the RT-clock of a node it is necessary to determine the temperature coefficients A_1 to A_3 . These coefficients are node specific and it is not possible to use a once determined set of coefficients for a set of nodes. However, it is possible to determine the temperature coefficients before node deployment and to use them for clock adjustments later.

The model given in Equation 2.14 was used to describe the measured clock changes in response to temperature variances of the Maxfor MTM-CM5000MSP RTC Clock (see Section 2.3.1). The determined temperature coefficients for each node are shown in Table 2.3. A comparison of the measured frequency-temperature relationship and the derived model are shown in Figure 2.18(a). It is important to note that this model only examines the effects on the frequency by temperature relative to the frequency observed at $25^\circ C$. To see how the results vary around the target frequency, a fixed offset (clock skew at $25^\circ C$) would need to be taken into account. For instance, node 13 has a clock skew of 60 ppm at $25^\circ C$ and this should be added to observe the drift in relation to the target frequency.

It can be seen in Figure 2.19 that the measured data fits the model with an average residual of 20.22PPM for all nodes across the measured temperature range, and a maximum residual of



(a) Maxfor MTM-CM5000MSP RT-Clock Model (b) STMicroelectronics MB950 Primary Clock Model

Figure 2.18: Model Fitting for RT-Clocks

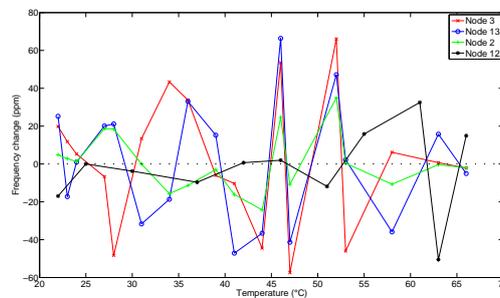


Figure 2.19: Maxfor RTC model Residuals

79PPM for nodes 3 and 13 between the temperature ranges of 46 °C to 53 °C.

We also applied the model to the primary systems clock of the STMicroelectronics MB950 as this primary system clock relies on a crystal. The measured frequency-temperature relationship and the derived model is shown in Figure 2.18(b). It be seen that these measured results again comply with the developed model however with a larger average residual of 129.44PPM and a maximum residual of 500PPM for node 2 at 23 °C.

2.4 Impact of temperature on sensing

One important function of a sensor device is measuring aspects of the environment. Many embedded devices measure properties such as temperature, sound and light which are used to form decisions on actuation. It is important that the values measured from the sensors provide the necessary accuracy to meet the application requirements. This makes the choice of the actual sensor used on the platform specific to the application. Whilst the environment may have an effect on such sensors, profiling all of these would be unrealistic.

Many of the environmental aspects sampled by sensors are analogue in nature and platforms make use of analog-to-digital converters (ADC) to represent sensor readings in a usable numerical format. Many MCU will contain ADC ports to which sensors can be connected to reduce

Node	A_1 [ppm/°C]	A_2 [ppm/°C ²]	A_3 [ppm/°C ³]
Node 12	1.2901	-0.39168	0.0071285
Node 13	1.1509	-0.23492	0.0017732
Node 2	0.93623	-0.12194	0.0012855
Node 3	1.6007	-0.27054	0.0041867

Table 2.3: Determined RT temperature coefficients for the evaluated Maxfor MTM-CM5000MSP nodes.

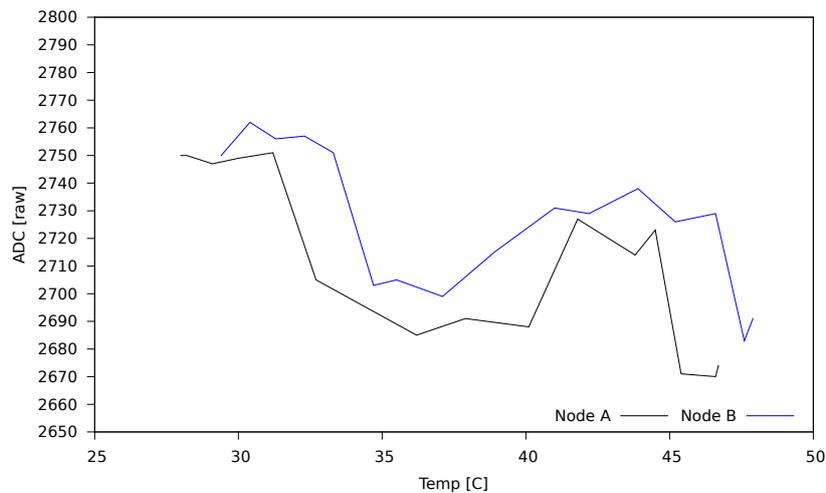


Figure 2.20: Maxfor ADC Sampling with Varying Temperature

the need for additional external circuitry. The ADC ports which the MCU provides must be capable to provide the necessary resolution and accuracy to support the application requirements. This section will look at how the accuracy of these ADC is effected by environment conditions, specifically temperature.

2.4.1 Experimental Results

To investigate how the accuracy of an ADC port is effected by changes in temperature, two Maxfor MTM-CM5000MSP motes were subjected to examinations. The ADC on the Maxfor mote has a resolution of 12 bits giving a maximum value of 4096. A constant stable voltage of 2 V provided by an Agilent N6705B-M1 DC power analyser [1] was connected to an ADC port which was sampled by the MCU. A Contiki-based application was created which continually looped and outputted the raw values sampled from the ADC. The temperature of the mote was increased by 20 °C using a heat lamp and dimmer module, described in previous experiments, over a 15 minute period. At each minute the ADC value and temperature was recorded.

Figure 2.20 presents the results from this experiment. A 2 V input on the ADC should produce a raw reading of 2730. Throughout all measured temperature ranges, the raw ADC values for both nodes were recorded to be within 60 units, or 1.5%, of this expected value.

The difference between the lowest and highest measurement for each node was 81 and 79 units, which equates to 2% variance range. Both nodes appear to exhibit similar trends but with no linear dependency between temperature and value.

2.4.2 Findings

Monitoring and reporting environmental aspects is the primary task of many embedded applications. It is important that the values measured are within accuracy bounds required by the application. This will depend on both the chosen sensor and ADC port if the sensor is an ADC. Only considering the ADC port, environmental temperature has been shown to effect the accuracy by up to 1.5% over a 20 °C range above room temperature. Whilst much research has gone into modelling temperature impact on crystals, with many documented models, there has been little focus on temperature effects on ADC sampling with no accepted published models. The results shown in Figure 2.20 suggest that the ADC values have a cubic relationship to temperature however more data is required provide a model to accurately approximate this relationship. Although we acknowledge that there is scope here to improve the accuracy of sensing, by modeling how temperature variances affect sensing accuracy, the focus of the RELYonIT project aims primarily to ensure reliable communication in the face of environmental conditions. The low impact of temperature variances against sensing should not have any impact upon inter-node communication. Therefore, given the limited time-frame of this project, the focus of the work will be to ensure reliable and deterministic communication. This work will be revisited should there be time later in this project.

2.5 Impact of temperature on energy supply and consumption

Many applications of embedded wireless systems operate on scarce energy supplied from either batteries or scavenged from the environment [32]. For such applications, changes in the energy supply or the rate that energy is consumed can have significant impacts on the lifetime of the application. It is important that we understand how changes in the environment affect both of these properties towards building more predictable and dependable platforms to support their applications.

Energy From Batteries. In wireless sensor networks nodes are often powered using batteries allowing applications to take advantage of the mobile nature of small embedded systems, and operating from batteries often facilitates lower installation costs as dedicated power outlets do not have to be built into the infrastructure where nodes are to be installed. Batteries are generally comprised of 3 main components, an anode and a cathode which are physically separated and submerged within an electrolyte solution. The electrolyte provides a medium to allow the flow of charge between the anode and the cathode.

It is possible for embedded systems to acquire energy from its environment from sources such as solar, wind, tidal etc. However when environmental energy sources are disrupted, nodes must depend upon other sources of energy (usually batteries) [32] to provide them with power. Energy harvesting therefore still relies upon the charging and discharging of batteries, which over a time period, will reach the cycle life of the battery and they will no longer be able to hold charge. For some battery technologies such as Ni-Cd batteries, the memory effect [34]

can be observed in which batteries gradually reduce the ability to hold charge through repeated charging after only partially discharging.

Battery profiles are determined, in part, to their chemical component. Common battery types include Nickel Cadmium (Ni-Cd), Nickel Metal-Hydride (NI-MH) and Lithium Ion, each of which offer different capacitance and respond differently under different discharge loads. As batteries deplete, two effects may alter the characteristics of the battery profile.

Rate capacity effects [26] affect the reachability of sites at the cathode. During low discharge rates the insoluble byproduct which affects sites of the cathode occur uniformly throughout the entire body of the cathode. During periods of high discharge, the compound affects only sites on the surface of the cathode, therefore rendering the inner sites of the material unreachable and hence reducing the usable capacitance of the battery.

Recovery effects [26] can recover some of the capacitance of a battery during idle periods where no current is drawn. During periods of large discharge, the rate of diffusion fails to keep up with the rate of which ions are consumed at the cathode. This results in a buildup on ions at the anode and a decrease of ions at the cathode. If the battery is allowed to idle for a period, charge recovery at the cathode occurs to decrease the imbalance, and as a result recovers some of the capacitance and lifespan of the battery.

Temperature Effects On Batteries. As temperature varies during the lifespan of a wireless sensor application, it has been shown in previous sections of this document that the frequency of all forms of platform clocks are affected. As energy consumption of embedded systems is related to the speed of which its components operate (amongst other factors), the rate of energy consumption and hence the discharge rate of the power supply of a node is likely to vary as temperature changes. It is unclear not only how this will directly affect the capacitance of a battery cell, but also how these variances will affect the rate capacity and recovery effects of these cells.

Node Energy Profiles. Knowing the exact rate that a platform consumes energy and how it fluctuates is important when selecting its energy supply. The energy supply must be capable of supporting the peak requirements of an application for a given platform, in addition to providing sufficient capacity to satisfy the lifespan requirements of the application. The failure to meet either of these requirements will lead to the platform being unable to support the application.

In an energy optimised system which enables components when necessary and duty-cycles them when not required, the rate of energy consumption will fluctuate as the application moves between states (generally off, sleeping, listening and sending). The peak energy rate as well as the total energy required to meet the application runtime can be predicted by analysing the power requirements.

2.5.1 Experimental Results

Profiling a platform will determine the necessary figures to enable the peak energy consumption rate, and total energy required, to be calculated. However, these measurements will only be accurate if either the environment has no effect on energy consumption, or the deployed environment is the same as the environment where the measurements took place. Therefore it

is important to understand how the environment effects the energy consumption of a platform. The following analysis will begin this investigation into how energy consumption of a platform is effected by its environment, by looking at energy consumption with varying temperature.

In the previous sections, variation in temperature has been shown to effect a platforms in numerous ways including timing, sensing and communication abilities, and it is likely that variations in temperature could affect the rate of energy consumption of a platform. To investigate the energy consumption rate of a number of platforms, the Maxfor MTM-CM5000MSP [29], Zolertia Z1 [50] and ST MB950 [38] were programmed to operate in a stable condition where the processor interrupts were disabled and a simple application ran in a continual, uninterrupted loop, and with external components disabled where possible. The platforms were then gradually heated by heat lamp used in previous experiements. Energy was supplied and the consumption rate measured using a calibrated Agilent N6705B-M1 DC power analyser [1].

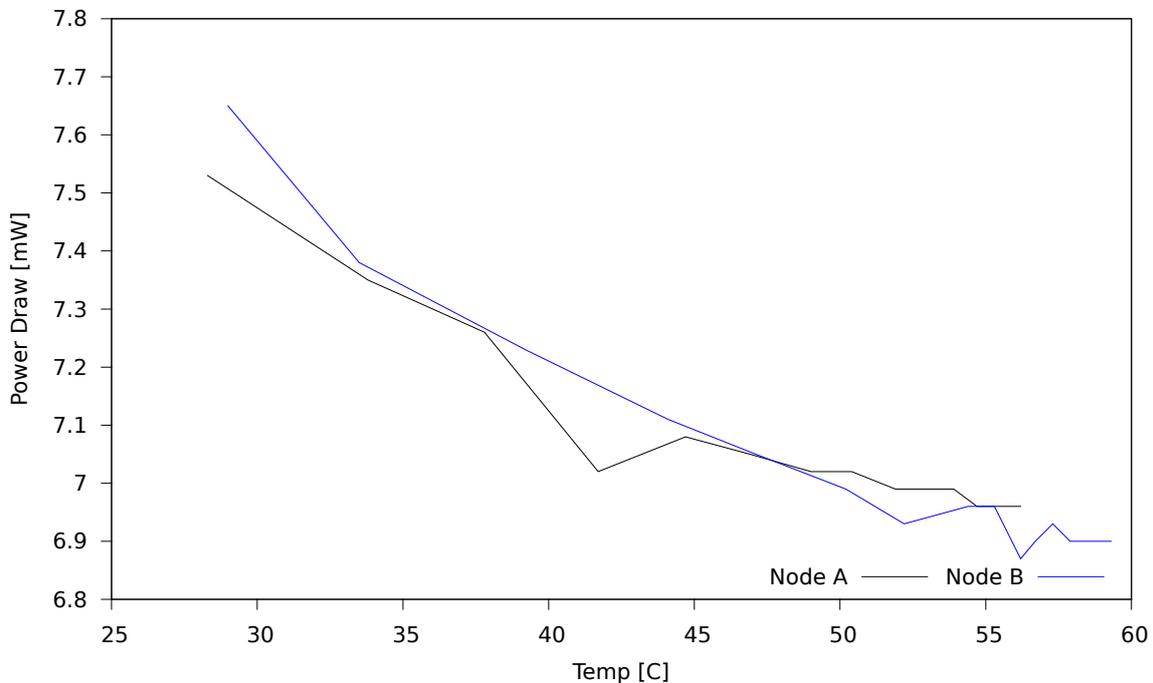


Figure 2.21: Maxfor Energy Consumption vs Temperature

The first platform examined was the Maxfor MTM-CM5000MSP mote. Figure 2.21 presents the results of this examination. Two nodes of this platform were tested and were found to follow approximately the same trend where, as the node is heated, the power consumption falls. At this point it is inconclusive if the relationship is linear as points on the plot become erratic. The power drawn of both nodes fell by an average of 9% over the 25°C temperature range tested. This fall in energy consumption closely correlates to the 13% reduction in processor clock rate that was observed for this platform.

A second platform examined was the Zolertia Z1. Figure 2.22 presents the findings of the investigation on this platform. It shows the opposite trend seen on the Maxfor mote, with energy consumption increasing with erratic behaviour at the both low and high temperatures

extremes with a difference of 3% between the minimum and maximum.

A final platform examined was the ST MB950 and Figure 2.23 presents its results. Again, similar to the Z1, the energy consumption appears to increase before sharply falling at 56 °C. The overall difference in the rate of consumption between minimum and maximum is 4%.

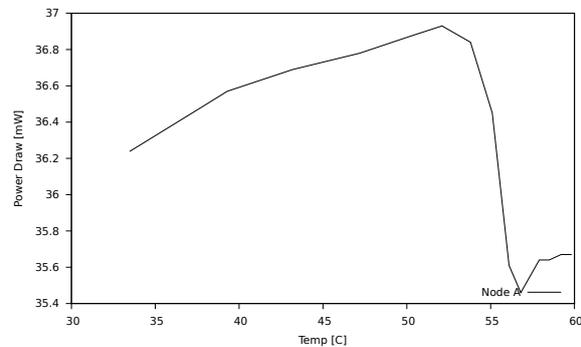
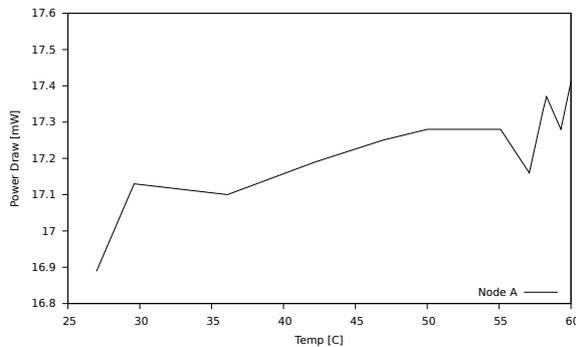


Figure 2.22: Zolertia Z1 Energy Consumption vs Temperature Figure 2.23: ST MB950 Energy Consumption vs Temperature

The results presented for each platform are very different and little to no commonalities can be seen. Creating a generic model to relate the energy consumption against temperature would be impossible without further analysis of the platforms.

2.5.2 Findings

To accurately represent the profiles of how power cells will discharge when providing power to embedded systems, there are a wide variety of variables to be considered, some which could be modeled such as the power profiles of each mode of operation for a given node, and the duration of time spent within these modes. However it is also unclear how the discharge effect and recovery effects will affect battery capacitance of relatively low power devices such as wireless sensor nodes.

According to battery manufacturer specifications [33] for the given temperature range of +20 °C - +70 °C, battery capacitance ranges from 2.35 Ah to 2.4 Ah when a constant current of 10 mA is being drawn, however, the capacitance between the upper and lower limits of the operational conditions for this battery yields a total variance of around 40%.

It can be found that the only major factor which environmental conditions affects in regards to power consumption/supply is the overall capacitance of a power source. The discharge rate should not be a factor for these low power nodes, as they operate within manufacturer specifications of the battery cells [33] for a modes of operation. It would be possible to model the power consumption of a node for all states, and we could model the environment and estimate how energy would be consumed over the lifespan of a WSN application, however we would not be able to estimate how the varying discharge rates would affect the chemical capacitance of batteries cell. It would therefore be beyond the scope of this project to attempt to model how varying discharge rates affect the chemical profiles of power cells.

3 Temperature - Environmental Models

As shown in the previous Chapter, temperature has a significant effect on the operation of sensornet platforms. There is henceforth a need to identify and model the environment's temperature profile. But what characteristics of the environmental temperature should we model? What temperature dynamics are important with regards to the operation of WSN? In this Chapter, we describe a set of metrics that we believe are important in modelling the environment. A thorough evaluation of these metrics is directly connected to the protocols to be modelled and evaluated (this will be done in subsequent deliverables).

3.1 Probabilistic Analysis

One of the main goals of the environmental model is to help in estimating the effect of temperature on sensornet platforms. Formally, if we let $f(t)$ be the model of one of the four "pillars" of a given platform (with t being a metric related to temperature), we would like a function $g : E \rightarrow T$ that maps the environment E to a meaningful metric for temperature t . The goal is hence to define a general function g to capture the environment.

Letting $g = pmf(S)$ be a function that outputs the probability mass function of a random variable S , the effect of temperature can be captured by the composite function $g \circ f(s)$. For example, for every temperature sample s in node i , we can apply the clock drift function defined in the previous chapter $f(s)$ and obtain the corresponding clock drift. Then, by applying $g \circ f(s)$ we obtain the probability mass function of clock drifts. This type of information would help the node in providing probabilistic guarantees, for example, by stating the likelihood that the clock drifts will go above a certain threshold $P(\text{clock drift} \geq th) = \sum_{t \geq e} g(e)$.

Considering the limited computing capabilities of nodes, the easiest way to implement the above explained method is to obtain the pmf of the temperature (through periodic sampling), and then, simply apply the function f to the domain of the pmf .

But not only the pmf of the temperature is important. When we evaluated the effect of temperature on the received signal strength, we noted that it is also important to capture how fast and abrupt the change in temperature can be. To obtain this type of information, we need to capture the first derivative of the temperature curve. Since our data is discrete, we cannot obtain a derivative, instead we calculate the slope of the curve at each point. Formally, denoting $S = \{s_0, s_1, \dots, s_n\}$ as consecutive samples of temperature during a day (or days), we obtain the slopes (first derivative) in the following way:

$$S' = \{s'_i = s_{i+1} - s_i, \text{ where } i = 1, \dots, n\} \quad (3.1)$$

Then, by applying the $pmf(.)$ function to S' we get a probabilistic distribution of the rate of change in temperature. Denoting $f(\Delta T)$ as the SNR function derived in the previous chapter (to capture the effect of temperature in the signal strength), $g \circ f(s')$ would denote the pmf of

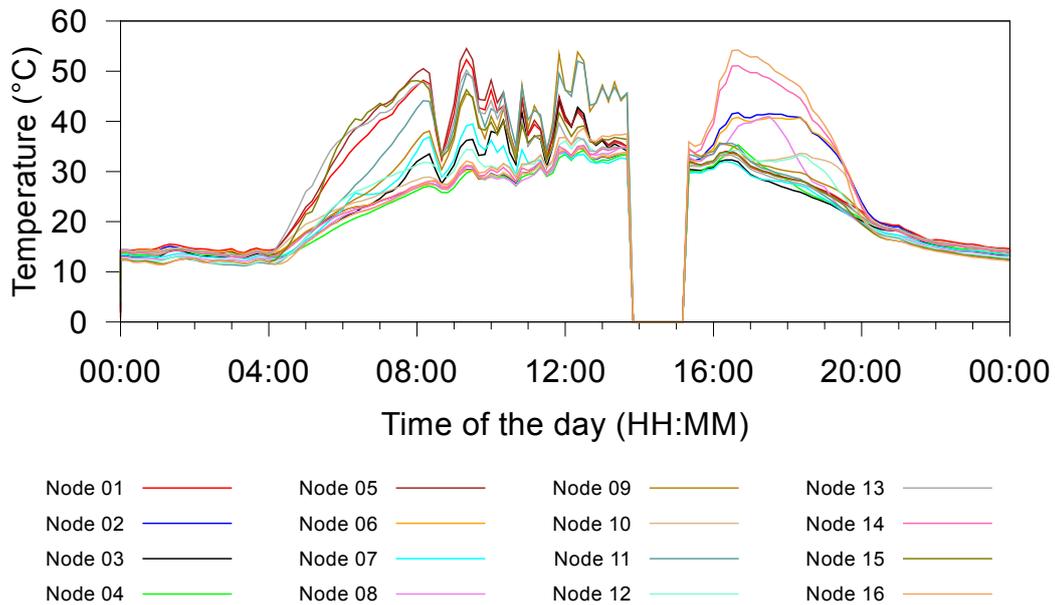


Figure 3.1: Temperature of 16 nodes during a day in an outdoor deployment. A temperature of 0 °C indicates that no data was collected during that time.

the change in signal strength. This type of information can estimate how exposed a node is to sudden changes in its link quality, $P(\Delta\text{RSS} \geq th) = \sum_{th \geq e} g(e)$.

3.2 Network Profile

One of the most important observations about the impact of temperature is that nodes are affected with different intensities. That is, given the same environment and the same time, the difference between the coldest and hottest nodes in the network can be of several tens of degrees. In terms of the network's operation, this implies that depending on their temperature profile, some nodes may be less able to transmit data, more likely to have clock drifts, more likely to have erroneous readings, or more likely to use more energy. In essence, this means that temperature enforces a high degree of heterogeneity into what would be otherwise a more homogeneous network. It is hence important to capture the degree of heterogeneity enforced by temperature.

Heterogeneity. Figure 3.1 depicts the temperature profiles of 16 nodes in the open field experiment performed in Sweden (see Section 2.1.1 for more details). The x-axis represents time and the y-axis represents the temperature of the nodes in °C. There is an important trend to highlight. During the night (in summer, Sweden has short nights), the temperature does not change much across nodes, i.e. temperature does not enforce heterogeneity. On the other hand, during the day we observe a high variation. To capture the degree of heterogeneity caused by

these variable trends, we use the Kullback-Leibler distance:

$$D(p||q) := \sum_{x \in X} p(x) \log \frac{p(x)}{q(x)}, \quad (3.2)$$

where $p(x)$ and $q(x)$ represent probability mass functions. Denoting $p_i(t)$ as the probability mass function of the temperature observed at node i (collected during a day), we define the "temperature distance" between nodes i and j by:

$$D(t_{i \leftrightarrow j}) = \frac{D(p_i||p_j) + D(p_j||p_i)}{2} \quad (3.3)$$

The greater the distance between nodes, the greater the heterogeneity caused by temperature. Denoting $G_N = (V, E)$ as the connectivity graph of network N , the heterogeneity imposed by temperature on the graph N is given by:

$$D_N = \sum_{(i,j) \in E} \frac{D(t_{i \leftrightarrow j})}{|E|} \quad (3.4)$$

The higher D_N , the higher the impact of temperature in the network. If, for example, we would divide the temperature profile in Figure 3.1 into day and night, we would clearly observe a much higher D_N for the day than for the night. For different environments, the D_N metric would provide a way to compare their relative exposure to temperature dynamics.

Global Max/Min Analysis. The heterogeneity metric derived above requires a very detailed knowledge of the network. For each node, we require the distribution of temperature and information about their neighbours. While in some scenarios the network may be able to gather this information during the first day(s) of operation (and then adapt accordingly), in other scenarios it may be necessary to give some minimum level of quality-of-service guarantees from the beginning. For the latter type of scenarios, we can estimate the maximum and minimum temperatures that any node would be exposed to, and then, plug in these values in our models for signal strength, timing, sensing, and energy consumption to obtain upper and lower bounds for the performance of the network.

3.3 Node Profile

In the previous section our focus was to get a macro view of the effect of temperature on the network. But protocols usually depend heavily on the individual characteristic of each node. For example, a routing or MAC protocol may prefer to avoid using "hot" nodes and use instead more reliable "cold" nodes. It is therefore important to obtain metrics that characterize the *individual temperature characteristics* of each node.

Hotness. Considering that the temperature observed by node i can be represented by a random variable T_i with probability mass function $p_i(t)$, we can use the coefficient of variation to capture the hotness of a node. Denoting μ_i and σ_i as the mean and variance of T_i , and \max_t as the maximal operational temperature of a node, we denote the hotness of node i by:

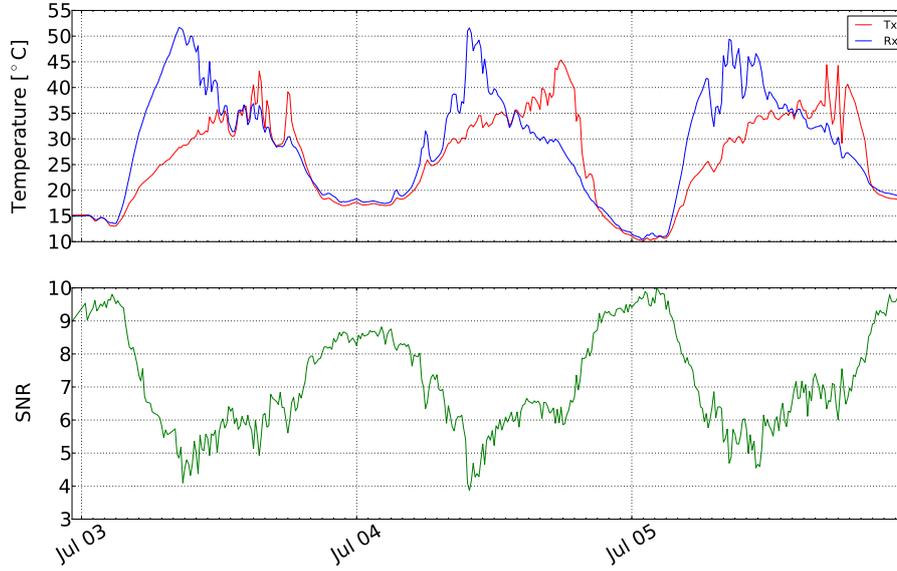


Figure 3.2: Periodic temperature pattern of a transmitter-receiver pair

$$C_i = \frac{\max_t - \mu_i}{\sigma_i} \quad (3.5)$$

The higher C_i , the colder the node and the lower its variance, which implies that the higher the C_i of a node, the more stable the node becomes (from the perspective of temperature).

Periodicity. Figure 3.2 depicts the temperature of two nodes for a few days. We can observe that both nodes follow different patterns (because they are exposed to sun light at different times during the day). This periodic behaviour can be important in the context of delay-tolerant applications. For example, if a node has a strong periodic behaviour, a delay-tolerant application could make the node sleep during periods of high temperature and wake it up during periods of lower temperatures.

To capture the periodicity of a node we use the concept of cross correlation. Denoting f_i and g_i as the time series of the temperature observed by node i at two different days, we first need to filter out the high correlation caused by the changes between day and night. The time series should include only the range between dawn (t_d) and sunset (t_s). In the experiments performed by Uppsala, t_d and t_s can be obtained through the information coming from the light sensors. The cross correlation between f_i and g_i is:

$$(f_i \star g_i)[n] := \sum_{m=-\infty}^{\infty} f_i^*[m]g_i[n+m] \quad (3.6)$$

Let us denote $t_{beg} = t_d^f - (t_s^g - t_d^g)$ and $t_{end} = t_s^f + (t_s^g - t_d^g)$. These times mark the beginning

and end of the cross-correlation calculation based on the times observed for dawn and sunset in two different days. For our purposes, we modify the basic cross-correlation definition to:

$$(f_i \star g_i)[n] = \sum_{m=t_{beg}}^{t_{end}} f_i[m]g_i[n+m] \quad (3.7)$$

Finally, the periodicity of node i is given by:

$$P_i = \max_{n \in [t_{beg}, t_{end}]} (f_i \star g_i)[n] \quad (3.8)$$

The higher P_i , the more predictable the impact of temperature becomes, which in turn makes it easier to adapt the performance of the network for periods where the impact of temperature is minimized.

Rate of change. An important goal of the RELYonIT consortium is to design protocols that are able to adapt to changes in the environment. To achieve this, it is important to guarantee that the time taken by a particular process is less than the rate of change of the environment. For example, if establishing the basic routing structure takes m minutes, it would be ideal to know that, in m minutes, the changes in the environment will not be high enough to prevent a process from reaching its steady state (since this could cause stability problems).

To capture the rate of change, we use the time series f_i from node i in the following way:

$$R_i = \max_t (f_i(t + \Delta t) - f_i(t)) \quad (3.9)$$

The lower R_i , the more stable the network is with regards to sudden changes in the environmental temperature. The value for Δt is determined by the protocol (or process) with the longest transient state in the network stack (usually these are protocols related to the Data Link Layer and above).

Local Max/Min Analysis. The metrics described thus far for individual nodes rely on storing and accessing a relatively high amount of information (probability density functions and time series). Even if the temperature samples are taken only every several tens of minutes, some sensornet nodes may not have the resources to do this. For highly resource-constrained nodes, we can obtain a node profile based on the max and min temperatures observed by that node. This would allow a simpler analysis of upper and lower bounds (similarly to what was explained for the Max/Min analysis in the Network Profile Section.)

3.4 Aggregating models

A desirable outcome of our modelling efforts is not only to *estimate* the performance of the network given some initial parameters, but also to *identify* the optimal (or near-optimal) parameters for a given deployment. Usually this type of problems can be solved via optimization methods such as linear or convex programming. Within this scope, the key challenge faced by RELYonIT is the large and complex nature of our platform models, which may increase the difficulty in formulating the constrained optimization problem.

To ameliorate this problem, we plan to utilize aggregation techniques to reduce the number of individual models. Due to the unique temperature profile of each node, the corresponding model will also be unique. We plan to cluster similar models into *cluster* models. For instance, neighbouring nodes having a similar but not identical temperature profile, will decide in a distributed manner to adopt a single cluster model. Neighbouring nodes can identify their similarity by utilizing the Kullback-Leibler distance. This is an initial hypothesis which will be further studied in the course of the project.

4 Radio Interference

An important factor that can significantly affect the performance of a wireless network is the presence of radio interference. The latter is caused by “neighbouring devices operating concurrently in the same frequency band, disturbing each other by emitting unwanted radio frequency signals that play havoc with the desired ones” [4].

Interference is a particularly severe problem for wireless sensor networks, as the presence of neighbouring devices transmitting at higher power largely increases the chances of hampering their low-power communications. If a low-power wireless sensor network is used in safety-critical scenarios such as health-care [14] and industrial control and automation systems [25], it needs to *guarantee* high packet delivery rates and limited delay bounds. Unreliable connections, as well as a reduced lifetime compared to the requirements, can not be tolerated.

In our study, we focus on external radio interference, i.e., the one caused by other appliances and radio technologies operating in the same frequency range of the network of interest¹, and devise models that capture how typical interference patterns vary over time and how they affect the operations of a given sensornet hardware platform.

In particular, we focus on the 2.4 GHz industrial, scientific and medical (ISM) band, a worldwide-available unlicensed portion of the radio spectrum that is shared with other radio technologies such as IEEE 802.11 (Wi-Fi), IEEE 802.15.1 (Bluetooth), and that is potentially threatened by the operations of ignition systems in engines or domestic appliances such as microwave ovens, cordless phones, baby monitors, game controllers, presenters, and video-capture devices [10], [35], [49].

4.1 Platform Model

In order to devise models that capture how interference affects the operations of wireless sensor networks, we start by answering the question: “what are the outcomes of radio interference on a given sensor node?”.

Packet loss. The primary outcome of radio interference is typically an increase in the packet loss rate, which may lead to high latencies and to an increase in the network traffic due to retransmissions. In the presence of a sufficiently strong interference signal, the receiver node is no longer able to discriminate the good signal from the interfering one. The receiver node can indeed reject any interference that is C_{Rej} weaker than the signal of interest, with C_{Rej} being the so called co-channel rejection capability of the transceiver (with unit dB).

¹Internal interference is instead the one generated by other wireless sensor nodes operating within the same network.

Co-channel rejection		-3		dB	Wanted signal @ -82 dBm. Undesired signal is an IEEE 802.15.4 modulated at the same frequency as the desired signal. Signal level for PER = 1%.
----------------------	--	----	--	----	---

Figure 4.1: Co-channel rejection as specified in the Texas Instruments CC2420 radio datasheet [41].

Any interfering signal stronger than that may result, depending on its duration and strength, in either a corrupted or a completely lost (i.e., not even detected) packet. The first case occurs when radio interference corrupts only some of the bits in a frame, leading to cyclic redundancy check (CRC) errors and a consequent dropped packet². In the vast majority of the cases, however, the radio does not even detect the presence of a frame.

We can hence build a simple model of packet loss due to interference as follows. Given a pair (A, B) of wireless sensor nodes in which A transmits a train of n packets $P_1 \dots P_n$ to B , a generic interfering signal will affect the reception of a packet P_i at node B as follows:

$$\begin{aligned}
 P_i &= \begin{cases} \text{received} & \text{if } (I_i - R_i) \leq C_{Rej} \\ \text{not received} & \text{if } (I_i - R_i) > C_{Rej} \end{cases} \\
 &= \begin{cases} \text{received} & \text{if } (R_i + C_{Rej}) \geq I_i \\ \text{not received} & \text{if } (R_i + C_{Rej}) < I_i \end{cases} \quad (4.1)
 \end{aligned}$$

where R_i is the received signal strength of packet P_i at node B , I_i is the signal strength of the interfering signal at node B during the reception of P_i , and C_{Rej} is the co-channel rejection of the radio transceiver used by B .

IEEE 802.15.4 radios typically specify their co-channel rejection threshold in their datasheet under “electrical specifications”. For example, Figure 4.1 shows the specifications of the co-channel rejection from the datasheet of the Texas Instruments CC2420 radio, i.e., $C_{Rej} = -3$ dB for a wanted signal of -82 dBm. A similar value can be found for the CC2530 radio, whereas the CC2400 datasheet specifies a co-channel rejection of -10 dB.

Assuming for example $R_0 = -79$ dBm and $C_{Rej} = -3$ dB, the presence of any interfering signal stronger than $I_0 = -82$ dBm will result in the (partial) loss of P_0 , which is equivalent of having a minimum Signal to Interference plus Noise Ratio (SINR) of $+3$ dB. Figure 4.2 shows the results of an experiment confirming this. We used a pair of two Maxfor MTM-CM5000MSP nodes exchanging packets at a rate of 128 packets/second and varied the temperature at the transmitter, so to generate a different output power (see Chapter 2) and added an interfering node [10] in proximity. As we can see in the figure, the interference I_i produced by the interferer

² In the presence of robust encoding schemes or forward error correction techniques (FEC), it may still be possible to reconstruct the packet [27], but with a significant energy expenditure. Liang et al. [27] have shown that in the presence of an IEEE 802.11b interferer, the ratio between corrupted packets and lost packets is typically less than $\frac{1}{5}$, whereas it is slightly higher in the presence of an IEEE 802.11g interferer.

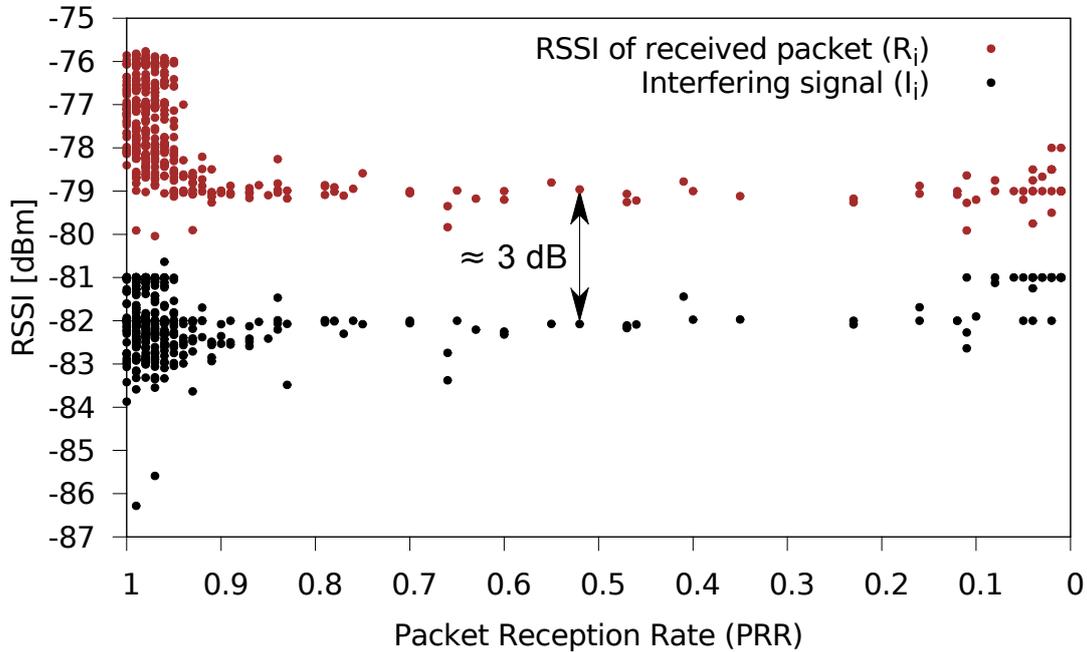


Figure 4.2: Co-channel rejection threshold found experimentally on the CC2420 transceiver, using Maxfor MTM-CM5000MSP nodes.

node is approximately -82 dBm. The weakest packets that can be received have $R_i \approx -79$ dBm: below this value, packets are not received. The difference between R_i and I_i is hence roughly 3 dB, which is as expected the inverse of the co-channel rejection ratio specified by the CC2420 datasheet.

Clear channel assessment. Interference may also affect the overall energy-efficiency and delay in the wireless network other than through loss of a packet. MAC protocols employing carrier sense multiple access with collision avoidance (CSMA/CA) typically sense the channel for ongoing transmissions and transmit packets only if the channel is found to be idle. The channel sensing procedure is typically called clear channel assessment (CCA), and whenever it detects the presence of an ongoing activity in the radio channel, the node defers or cancels the transmission, which may lead either to an unbounded latency or to the loss of a packet.

The clear channel assessment operation is typically based on the measured received signal strength compared against a programmable threshold T_{CCA} (in dBm units). In the popular Texas Instruments CC2420 platform, the CCA outcome can be read on a radio pin and the threshold level can be programmed in steps of 1 dB by programming the `RSSI.CCA_THR` register³.

³In the CC2420 radio, also a CCA hysteresis can be programmed in the `MDMCTRL0.CCA_HYST` control bits. The default CCA threshold is -77 dBm.

Therefore, the success or failure of the transmission of a packet P_i in the presence of a CSMA-CA protocol for a node can be simply modelled as:

$$P_i = \begin{cases} \text{transmitted} & \text{if } I \leq T_{CCA} \\ \text{not transmitted} & \text{if } I > T_{CCA} \end{cases} \quad (4.2)$$

where I is the signal strength of the interfering signal and T_{CCA} is the programmed CCA threshold.

Please notice that both Equation 4.1 and 4.2 do not depend on the actual strength of the interference signal, but only on whether it is above or below a given threshold. This can simplify our environmental models of interference, as discussed in the subsequent section.

4.2 Environmental Model

Most existing approaches have modelled interference on a device basis, i.e., they have modelled *individual sources of interference*. For example, Bianchi [7] has proposed an analytical model in which the Distributed Coordination Function (DCF) mode of 802.11 is modelled as a discrete Markov process where the back-off and retransmission mechanisms are represented as discrete states; whereas Taher et al. [40] have created an analytic model of the interference produced by microwave oven signals and studied its efficacy via simulation and experimental emulation.

In RELYonIT, we need to take a different approach to construct environmental models of radio interference. On the one hand, we do not have detailed knowledge of the deployment area that would allow us to determine model parameters for a given interference source prior deployment. On the other hand, the models need to be simple enough to be implemented on resource-constrained sensor motes, since the latter should be able to carry out a runtime assurance. In the remainder of this subsection, we describe a variety of interference models, ranging from very simple to more sophisticated ones, and try to keep them as generic as possible. The final choice on which model will be used will depend on the protocol models that will be considered in WP2.

4.2.1 Channel occupancy model

As we discussed in Section 4.1, to model the packet loss due to interference or the impact on CSMA-CA, one does not need to know the actual strength of the interference signal, but instead only whether it is above or below a given threshold. We therefore adopt the popular two-state semi-Markov channel occupancy model, in which, at a given time instant, a channel is defined as busy if any interfering signal is above a threshold R_{THR} and defined as idle otherwise [36]. The advantage of this simple model is that it can be easily used on constrained sensor nodes that are able to carry out energy detection, i.e., measure the received signal strength in absence of packet transmissions (we will refer to this as RSSI noise floor measurement [11] in the remainder of this document). Denoting x_i as the RSSI noise floor sampled by a node at a given time instant,

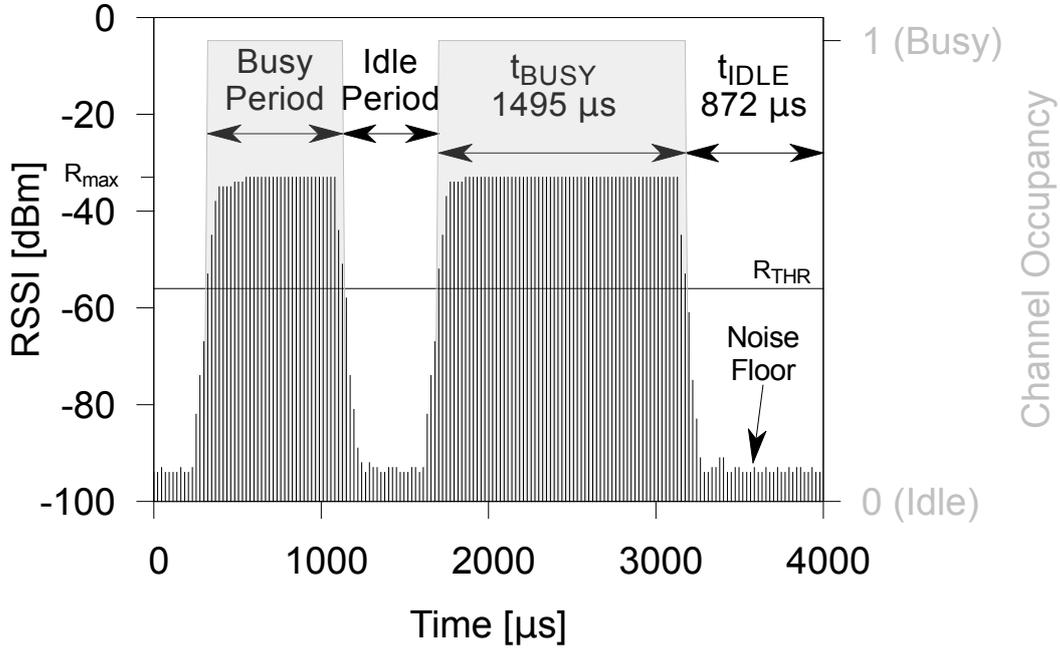


Figure 4.3: Example of interfering signal of strength R_{max} as recorded by a Maxfor MTM-CM5000MSP node performing a continuous RSSI noise floor measurements at a rate of 45.5 KHz. The shaded area shows the busy period of interference according to the two-state semi-Markov channel occupancy model. R_{THR} is set to -56 dBm.

the occupancy of the channel can be expressed as:

$$X_i = \begin{cases} \text{Busy (1)} & \text{if } x_i > R_{THR} \\ \text{Idle (0)} & \text{if } x_i \leq R_{THR} \end{cases} \quad (4.3)$$

with X_i being a binary number specifying a busy channel (1) or an idle channel (0), and R_{THR} being a user-specified threshold.

Figure 4.3 shows an example in which a Maxfor MTM-CM5000MSP node is performing continuous RSSI noise floor measurements at a rate of 45.5 KHz. Denoting $\{x_1, x_2, \dots, x_n\}$ as the sequence of consecutive RSSI noise floor measurements sampled at a rate of R Hz, and $\{X_1, X_2, \dots, X_n\}$ as the binary sequence of channel occupancy states computed according to Equation 4.3, one can derive an alternating sequence of idle and busy periods. The duration of idle and busy periods can be computed by knowing how many consecutive RSSI noise floor values are above or below R_{THR} . In the example from Figure 4.3, 35 and 65 consecutive RSSI noise floor measurements taken at $R = 45.5$ KHz were above R_{THR} , hence the two busy periods last 805 and 1495 μs , respectively.

The strength of the interference signals and the duration of idle and busy periods depend on

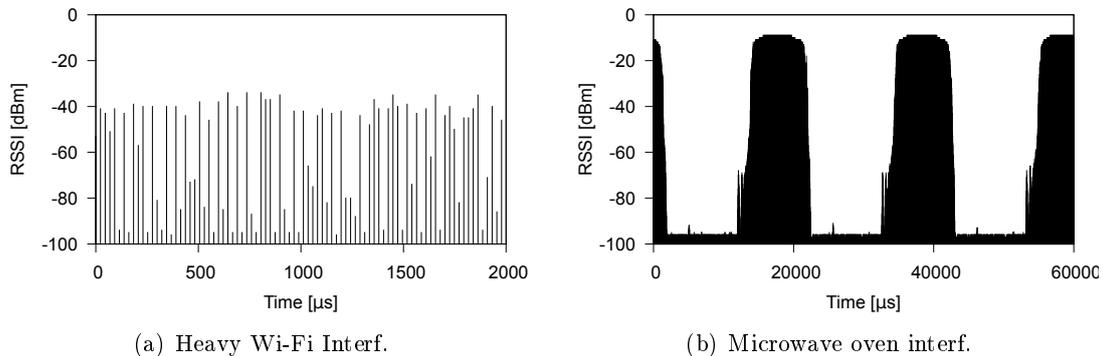


Figure 4.4: RSSI values measured using off-the-shelf wireless sensor nodes operating in the 2.4 GHz ISM band. Please notice the different scale of the x -axis [11].

the interfering source and on the protocol parameters. For example, the interference patterns generated by Wi-Fi transmissions depend on the number of active users and their activities, as well as on the traffic conditions in the backbone.

Wi-Fi transmissions are typically much stronger than sensornet transmissions, and can affect several IEEE 802.15.4 channels at the same time. Hauer et al. [18], [19] have shown that with a sufficiently high sampling rate, one can identify the short instants in which the radio medium is idle due to the Inter-Frame Spaces (IFS) between 802.11 b/g packets. Figure 4.4(a) shows the outcome of RSSI noise floor measurements at a rate of 45.5 KHz in the presence of heavy Wi-Fi interference caused by a file transfer: it is indeed possible to identify RSSI values matching the radio sensitivity threshold between consecutive Wi-Fi transmissions.

Figure 4.4(b) shows an example of the interference pattern caused by microwave ovens: high-power noise (≈ 60 dBm) is emitted in the 2.4 GHz frequency band in a very periodic fashion. The period mostly depends on the power grid frequency, but can also slightly vary depending on the oven model. Literature reports a power cycle of roughly 20 ms (at 50 Hz) or 16 ms (at 60 Hz) with an active period of at most 50% of the power cycle [10], [23].

4.2.2 Node Models

We now derive computationally lightweight interference models that can be implemented on resource-constrained sensor nodes.

Bounds on the duration of idle and busy periods. The simplest way to model interference is based on the minimum duration of an idle period min_{idle} and maximum duration of a busy period max_{busy} . Based on their knowledge, a CSMA-CA protocol could for example verify whether it is possible to guarantee to find an idle slot long enough to contain a packet of a given payload length, or to guarantee that a packet with a given payload can be transmitted within a certain deadline. Figure 4.5 shows a worst-case scenario analysis for the transmission of a packet using CSMA-CA. By knowing the time elapsed between two consecutive CCA checks t_{CCA} (assuming a constant back-off time), and the time t_{packet} necessary to

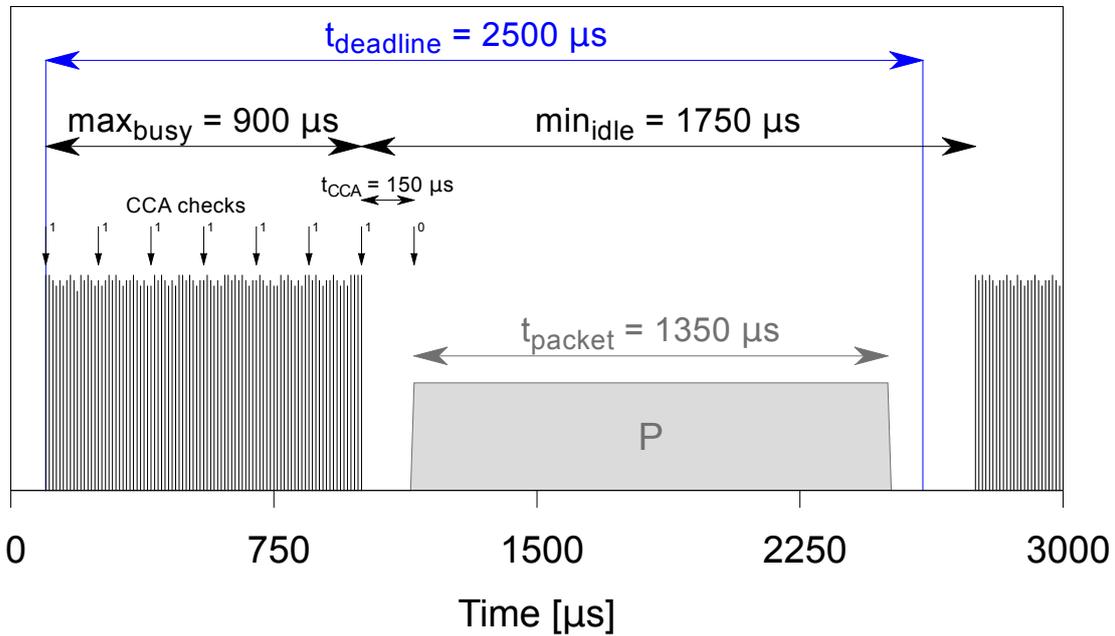


Figure 4.5: Worst-case scenario analysis for the transmission of a packet using CSMA-CA when using an interference model based on the minimum duration of an idle period min_{idle} and maximum duration of a busy period max_{busy} .

send the packet P over-the-air, one can verify if the transmission occurs within a certain time $t_{deadline}$. Assuming that the time elapsed between two consecutive CCA checks t_{CCA} is $150\mu s$, $min_{IDLE} = 1750\mu s$, and $max_{BUSY} = 900\mu s$, we obtain that a packet that needs to be transmitted within $t_{deadline} = 2500\mu s$ must take shorter than $t_{packet} = 1450\mu s$ to be transmitted over the air. The protocol can hence select a suitable payload length to guarantee the desired performance.

Cumulative distribution function of idle and busy periods. In principle, the longer the idle period, the higher the likelihood that a packet will be successfully received. For several protocol parameters, such as the CCA back-off time between consecutive busy channels, or the payload length, it is often important to know the actual distribution of idle and busy periods.

Figure 4.6 shows an example of the cumulative distribution function (CDF)⁴ of idle and busy periods measured by a Maxfor MTM-CM5000MSP node in the presence of a laptop continuously downloading a file from a nearby access point [11]. Please notice that to address the resource limitation of sensor nodes, the CDF retrieved by the Maxfor nodes was discretized. In such a scenario, the probability of having an idle period longer than 2 ms is smaller than 5%. This

⁴The cumulative distribution function of a real-valued random variable X is the function given by $F_X(x) = P(X \leq x)$, where the right-hand side represents the probability that the random variable X takes on a value less than or equal to x .

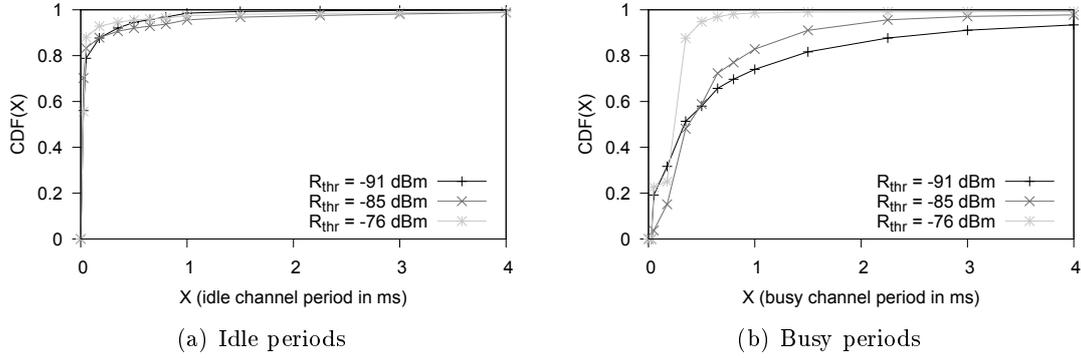


Figure 4.6: Cumulative distribution function (CDF) of idle and busy periods measured by a Maxfor MTM-CM5000MSP node in the presence of a laptop continuously downloading a file from a nearby access point [11].

is a hint for the protocol to use as short payloads as possible, and to avoid long CCA back-off times. On the contrary, an environment in which interference occurs in long bursts with large idle periods would probably call for long CCA back-off times in order to minimize the time in which the radio is on.

Denoting $p_i(i)$ as the probability density function of the idle periods formed by the interference pattern, a protocol could for example select the optimal payload length by computing the probability of encountering an idle period of length i :

$$p_{idle_period}(i) = \frac{ip_i(i)}{\sum_{i=1}^{\infty} ip_i(i)} \quad (4.4)$$

Similarly, denoting $p_b(i)$ as the probability density function of the busy periods formed by the interference pattern, a protocol could for example select the optimal back-off time for clear channel assessment by knowing the probability of selecting a busy period of length i :

$$p_{busy_period}(i) = \frac{ip_b(i)}{\sum_{i=1}^{\infty} ip_b(i)} \quad (4.5)$$

Modelling the interplay between idle and busy periods with a conditional CDF.

The knowledge of the cumulative distribution function of idle and busy periods is however not enough in case a protocol needs to exchange long sequences of packets, as in this case also the interplay between idle and busy periods needs to be known. Imagine a sensor node A that needs to exchange a sequence of messages with another node B . Node A sends a message P of duration t_P to node B that replies with an acknowledgement (ACK) message of duration t_{ACK} . In the presence of a cumulative distribution function such as the one in Figure 4.6(a), one would almost have a zero-probability of obtaining an idle slot long enough to contain both t_P and t_{ACK} . However, B requires a non-negligible time t_{load} to receive and extract P , analyse its validity, process its information, as well as to prepare the response frame and load it into the radio buffer. Hence, the presence of a busy period t_{busy} shorter than t_{load} after the transmission

of P would not have any impact on the reception of the ACK and the successful completion of the exchange.

Denoting $p_{idle}(i)$ as the probability of encountering an idle period of length i computed using Equation 4.4 and $p_{busy}(j < t_{load} \mid i)$ as the probability of obtaining a busy period j shorter than t_{load} after an idle period of length i , the probability of successfully completing the exchange p_{ex} is:

$$\begin{aligned}
 p_{ex} = & p_{idle}(i > t_P + t_{load} + t_{ACK}) + \\
 & \left(\sum_{i=t_P}^{t_P+t_{load}-\Delta t} p_{idle}(i) \times p_{busy}(t_P + t_{load} - i \mid prev_idle = i) \right) \times \\
 & p_{idle}(i > t_{ACK} \mid prev_busy = t_P + t_{load} - i)
 \end{aligned} \tag{4.6}$$

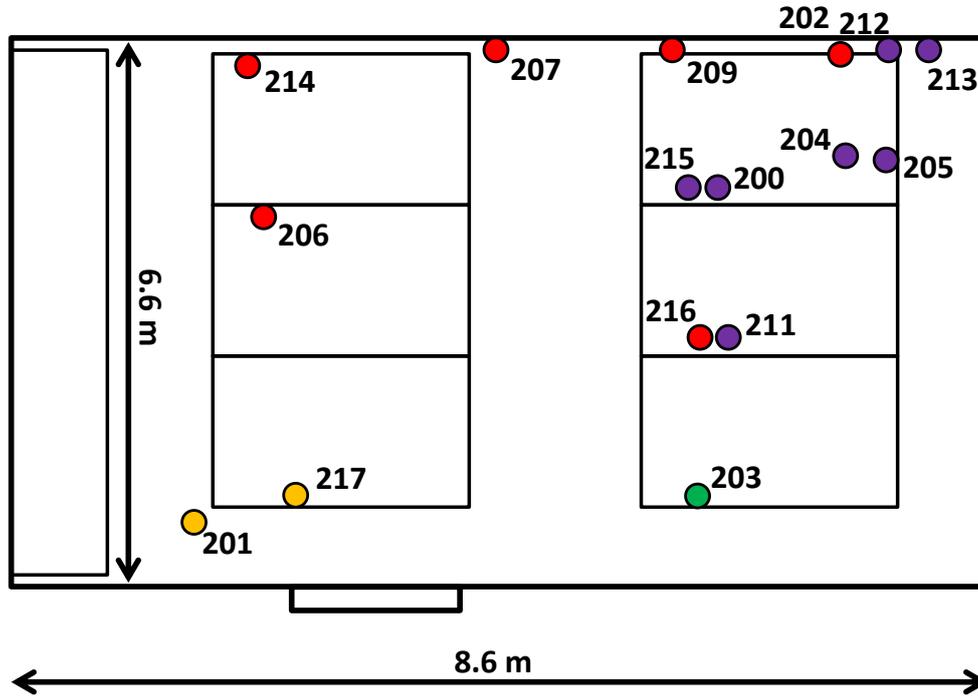
Node A can hence exploit the knowledge of the interplay between idle and busy periods to provide performance guarantees (see [11] for further details).

4.2.3 Network Model

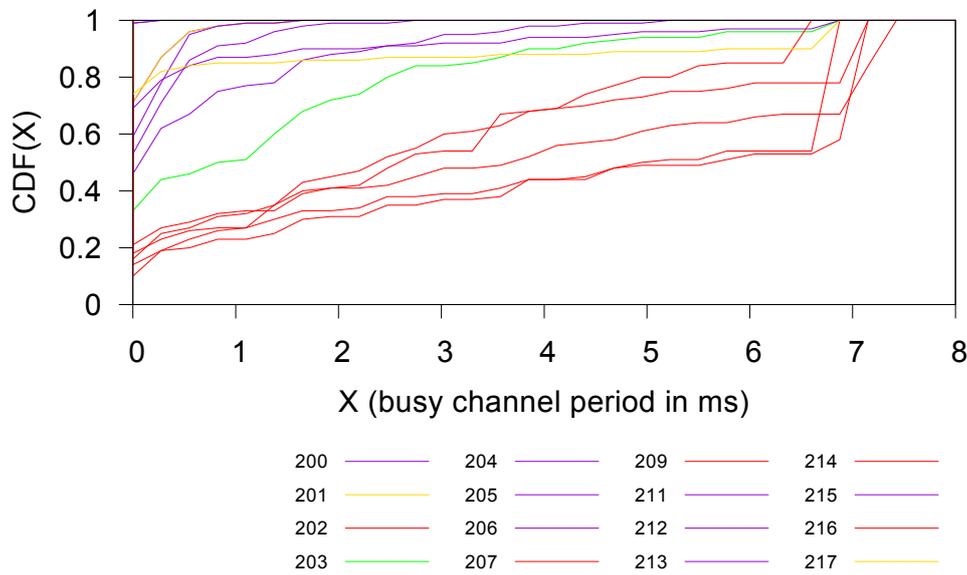
The models introduced in this section are typically computed on each node. External interference, however, usually affects a group of nodes in close proximity to each other in a similar fashion [20]. In order to avoid that every node in the network uses a different model and to minimize energy consumption, it would be important to select one model that is representative for a group of nodes or even for the whole network. Interference-aware protocols, for example, often need to adjust their parameters based on the interference measured by a group of nodes (e.g., a parent-children subset [20], [24]) and therefore need to select the most representative distribution, for example when selecting the most suitable channel [30].

Hence, we investigate the possibility of combining the measurements from a group of nodes and select an aggregate model that is representative for all of them. We first retrieve experimentally the duration of idle and busy periods among several Maxfor MTM-CM5000MSP nodes deployed inside a 56 m² office. The office is located in proximity of Wi-Fi access points and therefore rich of external interference, and the exact location of each node is shown in Figure 4.7(a). Each mote computes the distribution of idle and busy periods after collecting 50 million samples of RSSI noise floor values for a specific R_{THR} threshold. Figure 4.7(b) shows the CDF of busy periods measured on each node during night time with $R_{THR} = -89$ dBm: we can observe that external interference affects indeed nodes that are in proximity of each other in a similar way. In Figure 4.7(b) we can recognize two main trends highlighted in red and purple, and the corresponding nodes are highlighted using the same colours. As we can observe, they are indeed physically close to each other. In such scenario, the nodes belonging to red and purple can decide in a distributed manner to adopt a single cluster model, for example by identifying their similarity using the Kullback-Leibler distance as discussed in Chapter 3 and selecting the most representative distribution within a group⁵. This is an initial hypothesis which will be further studied in the course of the project.

⁵A cluster head can potentially also compute bounds within the aggregated CDF.



(a) Location of the nodes



(b) CDF of busy periods

Figure 4.7: Cumulative distribution function (CDF) of busy periods among several Maxfor MTM-CM5000MSP nodes located inside a 56 m² office. The colors identify the nodes in which similar interference patterns were measured.

5 Conclusions

The need to develop environmental and platform models arises from the fact that the performance of wireless sensor and actuator networks depends fundamentally on these two elements. Our efforts center around two environmental characteristics (temperature and interference) and four platform *pillars* (communication, timing, sensing, and energy consumption).

The topic of temperature has not received much attention from the research community, but in this report we have shown that it has a significant impact on the four *pillars* identified for the hardware platform. During the first eight months of work, most of our efforts have focused on temperature effects. On the other hand, the topic of interference has received significant attention from the community. For interference, our goal is to bridge the gap between analytical studies that rely on some unrealistic assumptions, and empirical evaluations that focus on specific platforms. Our general goal for the temperature and interference effects is to develop simple and generic models that are platform independent. These models will not only help in estimating the performance of sensornet deployments according to some initial set of parameters, but also in identifying optimal operating points.

As for the next steps, we have three main directions. First, to derive platform models for sensing and energy consumption. Second, to conclude the modelling for interference (for both, the environment and platforms). And third, to integrate the protocol models into this initial framework to fine-tune the models presented in this deliverable. It is also important to mention that our modelling efforts will be central to the tasks related to adjusting the parameters of the network in real-time.

Bibliography

- [1] *Agilent Technologies DC Power Analyzer, Model N6705*, Agilent Technologies, 2012.
- [2] *WiSMote datasheet - IPv6 platform for Wireless Sensor Networks R&D, Revision AS1006-BRF-201*, Arago Systems, feb 2012.
- [3] N. Baccour, A. Koubâa, L. Mottola, H. Youssef, M. Z. niga, C. Boano, and M. Alves, "Radio link quality estimation in wireless sensor networks: a survey," *ACM Transactions on Sensor Network (TOSN)*, vol. 8, no. 4, 2012.
- [4] N. Baccour, A. Koubâa, C. A. Boano, L. Mottola, H. Fotouhi, M. Alves, H. Youssef, M. A. Z. niga, D. Puccinelli, T. Voigt, K. Römer, and C. Noda, *Radio Link Quality Estimation in Low-Power Wireless Networks*, ser. SpringerBriefs in Electrical and Computer Engineering - Cooperating Objects. Springer, jul 2013.
- [5] K. Bannister, "Impacts of thermal reduction in transceiver performance on outdoor sensing networks," Master's thesis, Arizona State University, Phoenix, AZ, USA, 2009.
- [6] K. Bannister, G. Giorgetti, and S. K. S. Gupta, "Wireless sensor networking for hot applications: Effects of temperature on signal strength, data collection and localization," in *Proc. of the 5th Workshop on Embedded Networked Sensors (HotEmNets)*, 2008.
- [7] G. Bianchi, "Performance analysis of the IEEE 802.11 distributed coordination function," *IEEE Journal on selected areas in communications*, vol. 18, no. 3, pp. 535–547, 2000.
- [8] C. A. Boano, J. Brown, Z. He, U. Roedig, and T. Voigt, "Low-power radio communication in industrial outdoor deployments: The impact of weather conditions and ATEX-compliance," in *Proc. of the 1st International Conference on Sensor Networks Applications, Experimentation and Logistics (SENSAPPEAL)*, 2009.
- [9] C. A. Boano, J. Brown, N. Tsiftes, U. Roedig, and T. Voigt, "The Impact of Temperature on Outdoor Industrial Sensor Applications," *IEEE Transactions on Industrial Informatics*, vol. 6, no. 3, pp. 451–459, Aug. 2010.
- [10] C. A. Boano, T. Voigt, C. Noda, K. Römer, and M. A. Zúñiga, "Jamlab: Augmenting sensor network testbeds with realistic and controlled interference generation," in *Proceedings of the 10th IEEE International Conference on Information Processing in Sensor Networks (IPSN)*, apr 2011, pp. 175–186.
- [11] C. A. Boano, M. A. Zúñiga, K. Römer, and T. Voigt, "JAG: Reliable and predictable wireless agreement under external radio interference," in *Proceedings of the 33rd IEEE International Real-Time Systems Symposium (RTSS)*, dec 2012, pp. 315–326.

- [12] C. A. Boano, H. Wennerström, M. A. Zúñiga, J. Brown, C. Keppitiyagama, F. J. Oppermann, U. Roedig, L.-Å. Nordén, T. Voigt, and K. Römer, “Hot Packets: A systematic evaluation of the effect of temperature on low power wireless transceivers,” in *Under submission*, jun 2013.
- [13] Y. Chen and A. Terzis, “On the mechanisms and effects of calibrating RSSI measurements for 802.15.4 radios,” in *Proceedings of the 7th European Conference on Wireless Sensor Networks (EWSN)*, ser. LNCS 5970, feb 2010, pp. 272–288.
- [14] O. Chipara, C. Lu, T. C. Bailey, and G.-C. Roman, “Reliable clinical monitoring using wireless sensor networks: Experiences in a step-down hospital unit,” in *Proceedings of the 8th ACM International Conference on Embedded Networked Sensor Systems (SenSys)*, dec 2010, pp. 155–168.
- [15] *MICAz Datasheet, Revision A*, Crossbow Inc., 2004.
- [16] *TelosB datasheet, Revision B*, Crossbow Inc., may 2004.
- [17] S. Furber, *ARM System on Chip Architecture*. ADDISON WESLEY Publishing Company Incorporated, 2000. [Online]. Available: http://books.google.co.uk/books?id=J_Fu_YTVD9gC
- [18] J.-H. Hauer, V. Handziski, and A. Wolisz, “Experimental study of the impact of WLAN interference on IEEE 802.15.4 body area networks,” in *Proceedings of the 6th European Conference on Wireless Sensor Networks (EWSN)*, feb 2009, pp. 17–32.
- [19] J.-H. Hauer, A. Willig, and A. Wolisz, “Mitigating the effects of RF interference through RSSI-based error recovery,” in *Proceedings of the 7th European Conference on Wireless Sensor Networks (EWSN)*, feb 2010, pp. 224–239.
- [20] V. Iyer, M. Woehrle, and K. Langendoen, “Chryso: A multi-channel approach to mitigate external interference,” in *Proceedings of the 8th IEEE Communications Society Conference on Sensor, Mesh, and Ad Hoc Communications and Networks (SECON)*, jun 2011.
- [21] *Quartz Crystal Theory*, Jauch, 2007.
- [22] *JSJS Lightwave 300W Dimmer Datasheet*, JSJS, 2012.
- [23] A. Kamerman and N. Erkocevic, “Microwave oven interference on wireless LANs operating in the 2.4 ghz ISM band,” in *Proceedings of the 8th IEEE International Symposium on Personal, Indoor and Mobile Radio Communications (PIRMC)*, vol. 3, sep 1997, pp. 1221–1227.
- [24] B. Kerkez, T. Watteyne, M. Magliocco, S. Glaser, and K. Pister, “Feasibility analysis of controller design for adaptive channel hopping,” in *Proceedings of the 4th International Conference on Performance Evaluation Methodologies and Tools (VALUETOOLS)*, oct 2009.

- [25] L. Krishnamurthy, R. Adler, P. Buonadonna, J. Chhabra, M. Flanigan, N. Kushalnagar, L. Nachman, and M. Yarvis, "Design and deployment of industrial sensor networks: Experiences from a semiconductor plant and the north sea," in *Proceedings of the 3rd International Conference on Embedded Networked Sensor Systems (SenSys)*, nov 2005, pp. 64–75.
- [26] K. Lahiri, A. Raghunathan, S. Dey, and D. Panigrahi, "Battery-driven system design: a new frontier in low power design," in *Design Automation Conference, 2002. Proceedings of ASP-DAC 2002. 7th Asia and South Pacific and the 15th International Conference on VLSI Design. Proceedings.*, 2002, pp. 261–267.
- [27] C.-J. M. Liang, N. B. Priyantha, J. Liu, and A. Terzis, "Surviving wi-fi interference in low power zigbee networks," in *Proceedings of the 8th ACM International Conference on Embedded Networked Sensor Systems (SenSys)*, nov 2010, pp. 309–322.
- [28] S. Lin, J. Zhang, G. Zhou, L. Gu, T. He, and J. A. Stankovic, "ATPC: Adaptive transmission power control for wireless sensor networks," in *Proceedings of the 4th ACM International Conference on Embedded Networked Sensor Systems (SenSys)*, nov 2006, pp. 223–236.
- [29] *Maxfor MTM-CM5000MSP Datasheet*, Maxfor, 2010.
- [30] C. Noda, S. Prabh, M. Alves, C. A. Boano, and T. Voigt, "Quantifying the channel quality for interference-aware wireless sensor networks," *ACM SIGBED Review*, vol. 8, no. 4, pp. 43–48, nov 2011.
- [31] O. Rensfelt, F. Hermans, L.-A. Larzon, and P. Gunningberg, "Sensei-UU: a relocatable sensor network testbed," in *Proc. of the 5th ACM International Workshop on Wireless Network Testbeds, Experimental Evaluation and Characterization (WiNTECH)*, 2010, pp. 63–70.
- [32] S. Roundy, P. Wright, and J. Rabaey, *Energy scavenging for wireless sensor networks: with special focus on vibrations*. Kluwer Academic Publishers, 2004. [Online]. Available: <http://books.google.co.uk/books?id=PmmvRsfsUKcC>
- [33] *Primary lithium battery LS 14500 Datasheet*, Saft Batteries, 2008.
- [34] T. Sasaki, Y. Ukyo, and P. Novák, "Memory effect in a lithium-ion battery," *Nature materials*, 2013.
- [35] A. Sikora and V. F. Groza, "Coexistence of IEEE 802.15.4 with other systems in the 2.4 ghz-ISM-band," in *Proceedings of the IEEE Conference on Instrumentation and Measurement Technology (IMTC)*, may 2005, pp. 1786–1791.
- [36] L. Stabellini and J. Zander, "Energy-efficient detection of intermittent interference in wireless sensor networks," *International Journal of Sensor Networks (IJSNET)*, vol. 8, no. 1, pp. 27–40, jul 2010.
- [37] *STM32 32-bit MCUs*, STMicroelectronics, 2012.

- [38] *STM32W-RFCKIT RF control kit for STM32W108xx microcontrollers*, STMicroelectronics, 2012.
- [39] J. Sun and R. Cardell-Oliver, “An experimental evaluation of temporal characteristics of communication links in outdoor sensor networks,” in *Proc. of the 2th International ACM Workshop on Real-World Wireless Sensor Networks (REALWSN)*, 2006.
- [40] T. M. Taher, M. J. Misurac, J. L. LoCicero, and D. R. Ucci, “Microwave oven signal modeling,” in *Proceedings of the IEEE Wireless Communications and Networking Conference (WCNC)*, Apr. 2008, pp. 1235–1238.
- [41] *CC2420 datasheet - 2.4 GHz IEEE 802.15.4 / ZigBee-ready RF Transceiver, Revision SWRS041c*, Texas Instruments, feb 2013.
- [42] *CC2520 datasheet - 2.4 GHz IEEE 802.15.4 / ZigBee RF Transceiver, Revision SWRS068*, Texas Instruments, dec 2007.
- [43] *MSP430F1611 mixed signal microcontroller datasheet, Revision SLAS368G*, Texas Instruments, mar 2011.
- [44] *MSP430F2617 mixed signal microcontroller datasheet, Revision SLAS541K*, Texas Instruments, nov 2012.
- [45] *CC2530 datasheet - A True System-on-Chip Solution for 2.4-GHz IEEE 802.15.4 and ZigBee Applications, Revision SWRS081B*, Texas Instruments, feb 2011.
- [46] *CC2531 datasheet - A USB-Enabled System-On-Chip Solution for 2.4-GHz IEEE 802.15.4 and ZigBee Applications, Revision SWRS086A*, Texas Instruments, jun 2010.
- [47] J. Thelen, D. Goense, and K. Langendoen, “Radio wave propagation in potato fields,” in *Proc. of the 1st Workshop on Wireless Network Measurement (WinMee)*, 2005.
- [48] H. Wennerström, F. Hermans, O. Rensfelt, C. Rohner, and L.-A. Nordén, “A long-term study of correlations between meteorological conditions and 802.15.4 link performance,” in *Proc. of the 10th IEEE International Conference on Sensing, Communication, and Networking (SECON)*, 2013.
- [49] G. Zhou, J. A. Stankovic, and S. H. Son, “Crowded spectrum in wireless sensor networks,” in *Proceedings of the 3rd Workshop on Embedded Networked Sensors (EmNets)*, may 2006.
- [50] *ZOLERTIA Z1 Datasheet*, ZOLERTIA, 2010.