
University of Paderborn

Johannes Blanckenstein

Wireless Sensor Networks for Flight
Applications

Dissertation

submitted to the

Faculty of Electrical Engineering,
Computer Science, and Mathematics

in partial fulfillment of the requirements for the degree of

Doctor rerum naturalium (Dr. rer. nat.)

Munich, May 10, 2015

Referees:

Prof. Dr. Holger Karl, University of Paderborn

Prof. Dr. Falko Dressler, University of Paderborn

Additional committee members:

Prof. Dr. Stefan Böttcher, University of Paderborn

Dr. Simon Oberthür, University of Paderborn

Prof. Dr. Marco Platzner, University of Paderborn

Submitted May 10, 2015

Examination

Published

Munich, Germany

Abstract

The market forecast for aircraft manufacturers is very promising; the fleet of passenger aircraft will double. This will clearly generate a strong business for aircraft manufactures. But new competitors arise and, hence, rivalry is increasing. To succeed in this market situation, aircraft manufacturers have to build innovative aircraft to set themselves apart from competitors.

Most of the research effort is concentrated on developing lighter, more energy-efficient aircraft which reduce operational costs for airline operators and improve the flight experience for passengers to generate unique selling points. A very promising approach to accomplish these goals is to introduce wireless sensor networks for flight applications. Such wireless sensor networks can be very beneficial: they can help to reduce weight by saving cabling, they can improve workflows and, hence, reduce commissioning and operational costs, and they can enable new applications which were not feasible or even possible before.

In this work, flight applications are investigated to identify the challenges which arise when introducing such a wireless sensor network. Technologies and protocols are presented which aim to tackle these challenges. In particular, the most demanding prerequisites are energy efficiency, transmission reliability, scalability, synchronization, and localization.

Four of these demands will be addressed by three different protocols. First, a clock synchronization protocol is presented which uses a special hardware device—a wake-up receiver—to achieve synchronization in a very energy-efficient, reliable, and scalable way. Second, using this same technology a clustering protocol is presented which can reduce redundant transmissions. In doing so, it becomes possible to lower the mean energy consumption for hundreds of sensor nodes. Last, a custom-tailored medium access protocol is presented which utilizes spatial diversity to increase transmission reliability while keeping a very low power demand.

Kurzfassung

Die Prognosen der Marktentwicklung im Luftfahrtbereich sehen sehr positiv aus. In den kommenden 20 Jahren soll sich die Anzahl der Passagierflugzeuge verdoppeln, was sicherlich die Geschäfte im Luftfahrtbereich anregen wird. Jedoch bildet sich neue Konkurrenz in Asien, welche den Wettbewerb erhöhen wird. Um in dieser neuen Marktsituation weiterhin bestehen zu können, müssen Flugzeughersteller vermehrt innovative Flugzeugkonzepte entwickeln, mit welchen sie sich von ihren Konkurrenten absetzen können.

Die meisten Innovationen zielen auf eine Reduzierung des Gewichts und auf höhere Energieeffizienz von Flugzeugen ab. Ebenso steht eine Reduzierung der Inbetriebnahme- und Betriebskosten im Fokus. Ein weiterer wichtiger Aspekt ist das Generieren von Alleinstellungsmerkmalen, wie zum Beispiel das Schaffen eines außerordentlichen Flugerlebnisses, welches dazu beiträgt, dass Passagiere vorzugsweise in bestimmten Flugzeugtypen reisen wollen. Ein vielversprechend aussehender Ansatz diese Ziele zu erreichen, ist der Einsatz von drahtlosen Sensornetzen, um Luftfahrtanwendungen anzubinden. Der Einsatz so eines drahtlosen Sensornetzes kann in vielerlei Hinsicht Nutzen bringen. Verkabelung kann eingespart werden was große Gewichtsreduktionen mit sich bringt. Arbeitsabläufe können verbessert werden, wodurch Inbetriebnahme- und Betriebskosten reduziert werden können. Zusätzlich kann der Einsatz von drahtlosen Sendernetzen dazu beitragen, bisher nicht sinnvoll realisierbare Anwendungen einzuführen, beziehungsweise diese erst zu ermöglichen.

In dieser Arbeit werden typische Flugzeuganwendungen identifiziert, welche von dem Einsatz eines drahtlosen Sendernetzes profitieren können. Die Herausforderungen, die der Einsatz so eines drahtlosen Sensornetzes hervorruft, werden beleuchtet, als auch entsprechende Technologien und Protokolle vorgestellt, welche darauf abzielen, diesen Herausforderungen zu begegnen. Dabei wird das Hauptaugenmerk insbesondere auf folgende fünf Anforderungen gelegt: eine hohe Energieeffizienz, eine hohe Zuverlässigkeit des Kommunikationssystems, das Unterstützen einer sehr großen Anzahl an Kommunikationsteilnehmern, das Bewerkstelligen einer ausreichend genauen Zeitsynchronisation und das Bereitstellen einer Lokalisierungsfunktionalität.

Vier dieser fünf Anforderungen werden in dieser Arbeit adressiert. Hierfür werden drei Protokolle vorgestellt, welche verschiedenen Aspekte bei dem Betrieb von drahtlosen Sensornetzen abdecken. Zuerst wird ein Uhrensynchronisation-Protokoll vorgestellt, welches unter Verwendung eines speziellen Hardware-Bausteins — eines sogenannten Wake-up Receivers — dies besonders energieeffizient, zuverlässig und für zahlreiche Teilnehmer gleichzeitig bewerkstelligt. Zweitens wird ein Clustering-Protokoll vorgestellt, welches, unter Zuhilfenahme des gleichen Hardware-Bausteins, es ermöglicht, redundante Übertragungen einzusparen und somit den durchschnittlichen Energie-

verbrauch zahlreicher Sensorknoten stark zu reduzieren. Zuletzt wird ein speziell auf Luftfahrtanwendungen zugeschnittenes Medienzugriffsverfahren vorgestellt, welches die Übertragungszuverlässigkeit durch Ausnutzen von örtlicher Diversität erheblich verbessern kann; und dies ohne den Energieverbrauch entscheidend zu erhöhen, beziehungsweise unnötig stark zusätzliche Ressourcen zu verbrauchen.

Contents

1	Introduction	1
1.1	Motivation	1
1.2	Contributions and Thesis Structure	2
2	Background	7
2.1	Wireless Applications in Aircraft	7
2.1.1	Flight Test Instrumentation	7
2.1.2	Structural Health Monitoring	8
2.1.3	Wireless switch	10
2.1.4	Ambient Light Control	11
2.1.5	Application Requirements	12
2.2	Aircraft Radio Channel	13
2.3	Duty Cycling	16
2.4	Wake-up Receiver Technologies	18
2.4.1	Application Requirements	19
2.4.2	Technology	24
2.4.3	Comparison	26
2.4.4	Implementation Examples	28
2.4.5	Application Areas for Wake-up Receivers	32
2.4.6	Wake-up Receiver vs. duty-cycled Receiver	33
3	Energy-Efficient Clock Synchronization	37
3.1	Introduction	37
3.2	Background	37
3.2.1	Clock Oscillators	37
3.2.2	Periodic Synchronization	40
3.2.3	Delay Sources	41
3.2.4	Clock Synchronization Protocols	42
3.3	Bounded Clock Offset Synchronization	43
3.3.1	Keeping the Clock Offset Bounded	44
3.3.2	Jitter Analysis	45
3.4	Evaluation	49
3.4.1	Using a duty-cycled Transceiver	49
3.4.2	Using a Wake-up Receiver	51
3.5	Conclusion	53

4	Clustering using Wake-up Receivers	55
4.1	Introduction	55
4.2	Background	56
4.2.1	Sensor Readings	56
4.2.2	Mapping Measurements to Wake-up Messages	59
4.3	Protocol Design	59
4.3.1	Clustering	60
4.3.2	Monitoring	63
4.3.3	Protocol Extension	65
4.4	Evaluation	66
4.4.1	Clustering Phase	66
4.4.2	Overall Protocol Simulation	72
4.4.3	Verification by Hardware Implementation	78
4.5	Conclusions	79
5	Increasing Transmission Reliability	81
5.1	Introduction	81
5.2	Background	81
5.3	Protocol Description	82
5.3.1	Network Synchronization	83
5.4	Hardware Devices	84
5.5	Measurement Campaigns	85
5.5.1	Ariane 5	86
5.5.2	Airbus A330	92
5.5.3	Airbus A330 Mock-up	101
5.6	Conclusions	110
6	Conclusion & Outlook	111
	List of Figures	116
	List of Tables	117
	Nomenclature and Abbreviations	119
	Bibliography	136

1 Introduction

1.1 Motivation

Today around 32 million flights are conducted per year, transporting over 3 billion passengers and 50 million tons of freight. The Air Transport Action Group (ATAG) assessed the aviation industry's global economic impact at 2.4 trillion U.S. dollars annually, providing employment for nearly 60 million people [1]. On the downside flights are also responsible for over 700 million tones of carbon dioxide emissions, which amounts to over 2% of all human-induced emissions.

A market forecast [2] by Airbus predicted a compound annual growth rate of 4.7% for passenger traffic and of 4.5% for freight traffic in the next 20 years; the fleet of passenger aircraft will double. This high demand for new aircraft will generate a strong business for aircraft manufacturers. But despite such promising forecasts, the competition is growing stronger as well. Up to now, the market for aircraft with more than 100 seats is divided mostly between Airbus and Boeing, but new competitors from Asia are evolving. Additionally, the need to reduce carbon dioxide emissions grows stronger; aircraft with higher fuel efficiency are demanded.

To keep their current market share, aircraft manufacturers have to be able to offer a more “attractive” aircraft than their competitors. Hence, main goals for aircraft manufacturers are:

- to reduce the production / commissioning costs to increase the profit margin, respectively to decrease the unit cost;
- to reduce the operational costs for aircraft to make them more attractive for airline operators;
- to improve the flight experience and well-being of passengers to differentiate own aircraft from others.

To achieve these goals new technologies have to be introduced. The objectives of these technologies are:

- to improve aircraft-related production and commissioning workflows to decrease unit cost;
- to reduce the total weight of an aircraft to reduce fuel consumption;
- to improve the energy efficiency of all systems to reduce fuel consumption;

- to improve aircraft operation related workflows to reduce turn over time at airports;
- to improve flight experience to generate unique selling points.

The most effective way to reduce fuel consumption is to introduce new engines with greater efficiency. Aside from this, new materials like carbon fiber reinforced plastic (CFK) and new alloys can reduce the total weight of an aircraft. Weight reduction either reduces the fuel consumption or can allow to transport more passengers and cargo, which will increase the cost effectiveness per flight.

Weight is not only reduced by introducing new material but also by introducing new technologies. Cabling can be saved by introducing an aircraft-wide wireless sensor network (WSN). Several thousands of meters of cabling are installed in a typical aircraft. Many systems are connected by dedicated cables, not by a bus system. Replacing these cables by a wireless connection not only reduces weight but also simplifies and accelerates the production of an aircraft. Additionally, some systems only become possible when a WSN is introduced, because connecting the systems by cable is not feasible or even possible. For instance, integrating sensors into an aircraft structure only becomes possible when using a wireless connection. It also helps to introduce more flexible and reconfigurable systems. For example, without a cable connection, seats can be rearranged in minutes instead of days. The benefits of a WSN strongly depend on the characteristics of the individual aircraft subsystem, each subsystem / application has to be evaluated on its own.

In this thesis, the feasibility and the benefits of introducing wireless sensor networks for flight applications will be assessed. In particular, prerequisites to introduce such WSNs are identified and potential solutions are presented.

1.2 Contributions and Thesis Structure

Five major requirements for WSNs for flight applications are identified in Section 2.1.5: energy efficiency, transmission reliability, scalability, synchronization, and localization. Out of these five requirements, four will be addressed in this work; either by new protocols or by using a special hardware device—a wake-up receiver (WUR) as introduced in Section 2.4. The prerequisite localization will not be addressed in this thesis.

The PERT chart in Figure 1.1 depicts the thesis structure and chapter dependencies. In Chapter 1 the motivation for this work is presented. In Chapter 2 the common background information for the subsequent chapters is presented. Additionally, in Chapter 2 WURs are introduced and compared to state-of-the-art technologies. This information about WURs is only necessary for Chapter 3 and 4; the protocol presented in Chapter 5 does not depend on WURs. The main contributions of this thesis are presented in the subsequent Chapters 3, 4, and 5. Therein different protocols are presented that are

necessary when introducing WSNs for flight applications. The thesis is then concluded in Chapter 6.

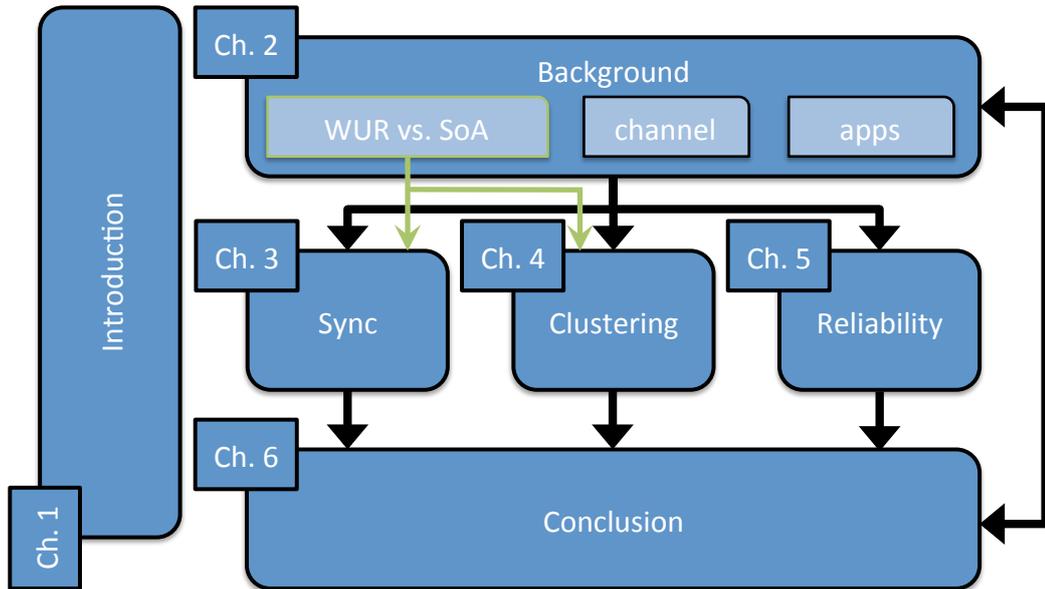


Figure 1.1: PERT chart of the thesis structure.

The following list provides an overview of the topic of each chapter and outlines its main contributions. The addressed prerequisites per chapter and the related earlier publications are presented, as well.

Chapter 1: Introduction The first chapter outlines the motivation for this thesis, presents its structure, and states the main contributions.

Chapter 2: Background Background information about main aspects for WSNs for flight applications is given. First, four typical aircraft applications are presented, which benefit from using a WSN, and their requirements are outlined. Then, a short overview of the aircraft radio channel is given. Last, an introduction to duty cycling and a summary of low-power receivers and their technology is presented. This summary about WURs is already published in the following magazine paper:

- Johannes Blanckenstein, Jirka Klaue, and Holger Karl. A Survey of Low-Power Transceivers and their Applications. Accepted (IEEE Circuits and Systems Magazine), pages 1–10, 2015

Chapter 3: Energy-Efficient Clock Synchronization Many flight applications rely on synchronized local times. In this chapter, an energy-efficient way to create a common time base is presented: wake-up receivers are utilized to synchronize clocks in an energy efficient way. Hence, the prerequisites energy efficiency, reliability, scalability, and

synchronization are addressed. Parts of this chapter are already published in the following paper:

- Johannes Blanckenstein and Holger Karl. Energy-Efficient Clock Synchronization using Wake-up Receivers. In 22nd International Conference on Software, Telecommunications and Computer Networks (SoftCOM 2014), pages 1–6, Split, September 2014

Chapter 4: Clustering using Wake-up Receivers Several hundreds of sensors are distributed throughout an aircraft to obtain various measurements. To reduce the power consumption of these sensor nodes an energy-efficient clustering protocol is presented that avoids redundant transmissions. This clustering protocol utilizes wake-up receivers and is aligned for the aircraft environment. Hence, the prerequisites energy efficiency, and scalability are addressed. Most of this chapter is already published in the following papers:

- Johannes Blanckenstein, Jirka Klaue, and Holger Karl. Energy Efficient Clustering using a Wake-up Receiver. In 18th European Wireless Conference (EW 2012), pages 1–8, Poznan, April 2012. doi:ISBN978-3-8008-3426-9
- Johannes Blanckenstein, Jirka Klaue, and Holger Karl. Reducing Redundant Transmissions in WSNs by Clustering using Wake-up Receivers. Submitted (Ad Hoc Networks), pages 1–13, 2015

Chapter 5: Increasing Transmission Reliability Many flight applications are safety-critical and thus demand a high transmission reliability. Because of scarce bandwidth resources and high network load, retransmissions are not the best solution for increasing transmission reliability. In this chapter a protocol is presented that introduces spatial diversity by using multiple access points per wireless cell. It is capable of increasing transmission reliability while increasing the resource demands only slightly. This protocol is evaluated in real-world flight environments: in the upper stage of a space launcher Ariane 5, in an aircraft Airbus A330, and in an aircraft cabin mockup of an A330. The prerequisites energy efficiency, reliability, and scalability are addressed. Parts of this chapter are already published in the following papers:

- Johannes Blanckenstein, Javier Garcia-Jimenez, and Jirka Klaue. A Scalable Redundant TDMA Protocol for High-density WSNs Inside an Aircraft. In Koen Langendoen, Wen Hu, Federico Ferrari, Marco Zimmerling, and Luca Mottola, editors, Real-World Wireless Sensor Networks, volume 281 of Lecture Notes in Electrical Engineering, pages 165–177. Springer International Publishing, 2014. doi:10.1007/978-3-319-03071-5_18

- Johannes Blanckenstein, Cristina Nardin, Jirka Klaue, and Holger Karl. Error Characterization of multi-Access Point WSNs in an Aircraft Cabin. In 1st IEEE ICC 2015 Workshop on Dependable Vehicular Communications (DVC), pages 1–6, London, June 2015

Chapter 6: Conclusion & Outlook The last chapter concludes this thesis and presents an outlook about future technology developments for aircraft WSNs.

2 Background

In this chapter, first, four typical aircraft applications are presented and their requirements are outlined. Second, a short overview of the aircraft radio channel is given. Third, the concept of duty cycling is presented, as well as a summary about low-power receivers. Finally, both are compared to each other.

2.1 Wireless Applications in Aircraft

Applications which are using wireless sensors are expected to be very beneficial in future aircraft. Their operational areas are manifold and concern many different aircraft systems. Four different applications are highlighted in the following, which represent the most promising applications in the near future: Flight Test Instrumentation (FTI), Structural Health Monitoring (SHM), Wireless Switch (WS), and Ambient Light Control (ALC).

2.1.1 Flight Test Instrumentation

After an aircraft is manufactured the functionality of all the systems has to be checked—the phase of commissioning starts. One part of the commissioning is the flight test, where the behavior of the aircraft is monitored during its first flight.

In order to monitor stress, temperature, humidity, acceleration of the aircraft structure during the flight a measuring system—the Flight Test Instrumentation—has to be installed throughout the whole aircraft. Several Engineers are occupied for approximately one week to install the FTI. Then, a specific flight test is performed, typically within one day. Afterward, again several engineers are occupied for several days to deinstall the FTI. This is done several times for each specific test to perform. This huge effort to install and deinstall the FTI is evoked by the complexity of the measuring system and the great amount of sensors. Several thousands of meters of cable have to be installed at proper locations, which not only amounts to weights of several thousands of kilos but more importantly is also error-prone. Such an FTI is depicted in Figure 2.1.

To install the sensors at the correct locations and to connect them with the correct port at the correct measuring racks is highly complex and, thus, takes a lot of time.

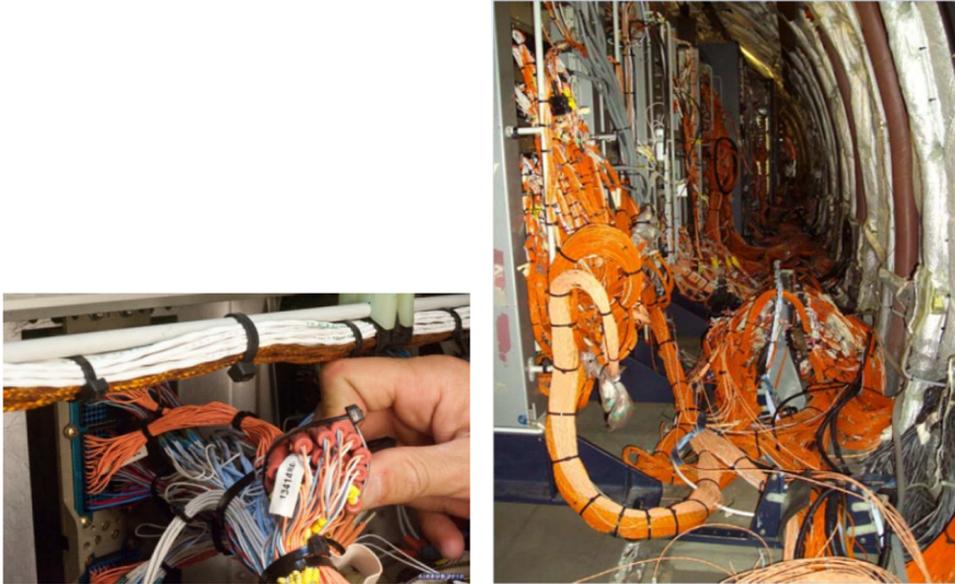


Figure 2.1: Flight Test Instrumentation (From Airbus S.A.S.).

The idea to improve the flight test process is to connect the sensors wirelessly to the measuring racks. The complexity of the cabling process and hence its error-proneness could be radically reduced. Then it should be possible to install and deinstall the FTI in only a few days. The commissioning time can be reduced and, consequently, the costs of the aircraft can be lowered / the profit margin for the aircraft manufacturer can be increased.

Of course, such a WSN has its own challenges. It has to be capable of supporting several thousands of sensor nodes. The provided data rate for the sensor nodes has to be high enough and the sensor nodes have to be synchronized with a high enough accuracy to align the measurements correctly. Additionally, the sensor nodes have to be self-sustained because a power supply line would counteract the improvement by the wireless data connection. Hence, the nodes have to have enough energy to support an operation time at least as long as the flight test duration including the installation phase, either by having a large enough battery or by using energy harvesting. Last, a sensor node, respectively the WSN, has to be capable of auto-localization to verify the correct position of the sensors.

2.1.2 Structural Health Monitoring

Another potential application for WSNs is Structural Health Monitoring. Up to now, the structural health of an aircraft is checked in fixed intervals. The check intervals are aircraft type-specific and are categorized in A, B, C, and D checks. Whereas A and B checks occur in relatively short intervals (each two months for the A check and about each three to four months for the B

check), C and D checks have much longer intervals (15 to 18 month for the C check and around six to ten years for the D check). A and B checks are minor maintenance checks, C and D checks are major maintenance checks where the aircraft is more or less been taken apart. It can even be necessary to remove the paint to inspect the metal skin of the fuselage directly. Such maintenance checks clearly involve a tremendous effort and expenditure. A D check can demand up to 50,000 man hours of work.

To decrease the immense maintenance costs for aircraft operators, it is proposed not to have fixed-scheduled maintenance intervals but to perform maintenance-on-demand instead. Maintenance-on-demand means to only schedule a maintenance when it is needed. To achieve this, the structural health of an aircraft fuselage has to be monitored to recognize when a maintenance is necessary. A lot of sensor nodes have to be integrated into the aircraft fuselage to provide such a monitoring system (as illustrated in Figure 2.2).



Figure 2.2: Structural health monitoring.

Besides the challenges of integrating sensors into the fuselage at all, it is obvious that it is not desirable to drill holes through the fuselage to connect these integrated sensors by wire, regardless their usage for power supply or data communication. Holes are not only prone to act as heat bridges but can also weaken the structural integrity of the fuselage. Hence, the sensor nodes have to be self-sustained and wirelessly connected. Similar challenges as for FTI occur for SHM. A lot of sensor nodes are needed to monitor the whole fuselage, the sensor nodes have to be synchronized to align their measurements to each other in time, and the sensor nodes have to be energy self-sufficient. The big difference of SHM to FTI is that SHM nodes are installed permanently while FTI nodes are only installed for around one week. Hence, SHM nodes typically have to use energy harvesting to fulfill their energy demand over their entire lifetime. But SHM nodes typically

require only a lower data rate because most of the effects are slow (like corrosion) or only have to be monitored if certain thresholds are exceeded or events occurred. Additionally, because a maintenance procedure has to be scheduled and cannot be done during flight, it is sufficient to transmit the SHM data once a day to a central entity and has not to be done in real time.

2.1.3 Wireless switch

The status of many aircraft systems has to be monitored; usually in a binary fashion, e.g., if a door is opened or closed, if a switch is turned on or off, et cetera. Up to now, either dedicated wires are run directly to the sensors or the sensors are not monitored at all because costs are too high. One example of such a monitoring system is the torque limiter indicator. Under certain conditions it is possible that the gear boxes to drive the flaps and slats of an aircraft wing can freeze. To prohibit damage on the gear boxes torque limiters are installed. If a torque limiter is deployed it will release a small pin to indicate its deployed status. This indication is purely mechanical, hence, if a malfunction of the flap / slat systems is determined service technicians have to inspect all torque limiters to discover which gear box has to be defrosted. Which means using a ladder to reach all torque limiters on all flaps / slats. Clearly a time consuming operation. If the status of such a torque limiter indicator were monitored, the service technician could go directly to the frozen gear box to defrost it. Unnecessary down-time could be reduced. Connecting the torque limiters by wire is too costly, hence, a wireless system which monitors them would be a great benefit for aircraft operators. In Figure 2.3 such a wireless torque limiter monitoring system is depicted.

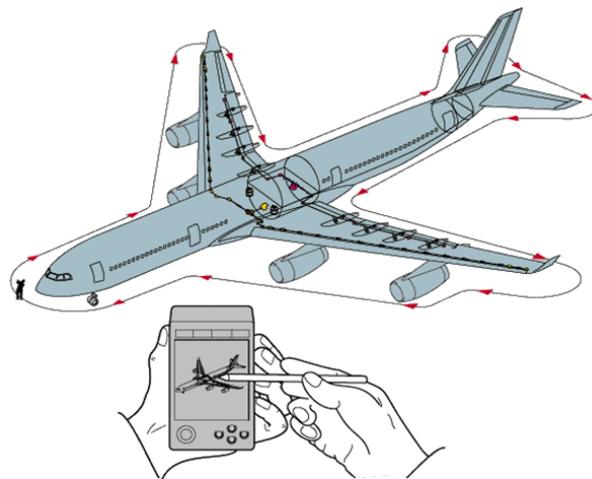


Figure 2.3: Wireless switch for torque limiter indicator.

The torque limiter indicator is only one of many applications where it is

beneficial to monitor, respectively control, a switch wirelessly. All the systems have in common that they are event-driven and most of the time only have to be used before start or after landing. But, as for the other applications as well, there are a lot of these systems, which can be challenging for a WSN to support all of them.

2.1.4 Ambient Light Control

One major selling point for aircraft manufacturers is—besides the efficiency of an aircraft—the flight impression and well-being of passengers. A type of aircraft can be more attractive for operators than others if it has a unique selling point, as for example a spectacular ambient light, which will motivate passengers to preferably fly with it.

Ambient light is not only consciously observed when for example a sunrise propagating through the cabin is simulated, but is also subconsciously important to affect the mood and behavior of passengers. For example, a quiet atmosphere can be created to settle down passengers and to help and prepare them for the night-time flight. Such an ambient light installation can be observed in Figure 2.4 where a day-time flight ambience is created.

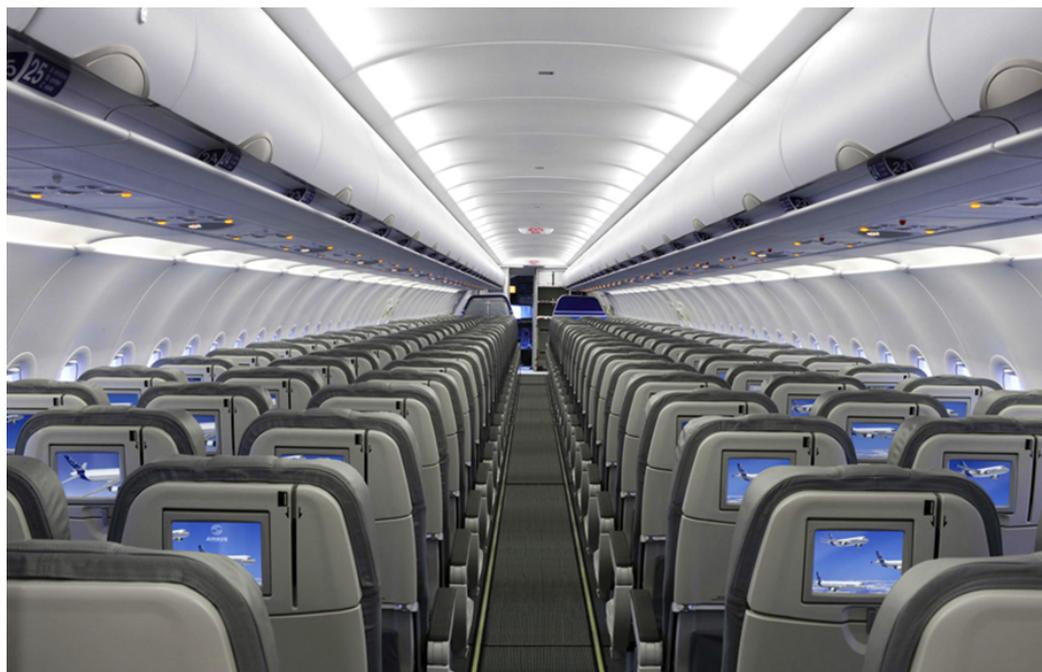


Figure 2.4: Ambient Light Control (From Airbus S.A.S.).

To create impressive ambient light, a lot of light installations are necessary. Additionally, the light installations cannot simply be turned on or off but have to be regulated in a more complex way. Up to now, each light installation is connected by a dedicated cable to a cabin control backbone network. This

dedicated cable is also used to identify, respectively assign, the location of the light installation. Hence, the assembly of light installations is a complex and time-consuming process. If the light installations were connected wirelessly, it would not only be possible to reduce the assembly effort but also to save a lot of weight in cabling. Clearly, a wireless node, respectively the WSN, has to be capable of auto-localization to support an easy assembly. Additionally, the wireless nodes have to have a common time basis to coordinate complex ambient light scenarios, hence, their clocks have to be synchronized. In contrast to other wireless applications, the power constraints of the ambient light wireless nodes are simple because the light installation will be connected to a power supply line in any case.

2.1.5 Application Requirements

Several challenges arise when introducing a WSN in an aircraft environment. These four groups of applications stated before cover most of the different aspects of aircraft WSNs. To compare their challenges and commonalities different key aspects are investigated in Table 2.1.

Table 2.1: Wireless application requirements

	FTI	SHM	WS	ALC
energy-efficiency demand	medium	high	high	low
transmission reliability	high	medium	high	medium
life span before access	1 week	10 to 25 years	6 to 10 years	5 years
number of entities	~ 1000	~ 500	~ 100	~ 100
synchronization	necessary	necessary	N/A	necessary
localization	necessary	necessary	improvement	necessary
type of occurrence	periodic	periodic	event-based	event-based

Most of the aircraft applications demand a very long lifetime, hence, meeting the power demands is one of the major goals of an aircraft WSN. This

demand can be addressed by different means: reducing the power consumption by using special hardware devices, increasing energy efficiency by introducing specially designed protocols, and by using energy harvesting, if applicable. Additionally, aircraft WSNs are very dense networks with more than several hundreds of nodes in a relatively small area. This density can be a great challenge regarding the used Medium Access Control (MAC) scheme. For example, carrier sense multiple access (CSMA) usually used in WSNs will suffer a lot of packet collisions and, hence, lower transmission reliability and throughput immensely. More suitable MAC schemes are scheduled, contention free time division multiple access (TDMA) variants, which are capable of preventing packet collisions.

Such a TDMA scheme has an extra benefit by being deterministic: all aircraft systems have to be certified by the International Civil Aviation Organisation (ICAO) or their local representatives. Several aspects have to be addressed to become certified. Amongst others, it has to be proven that a system fulfills special safety requirements. The determinism of a MAC scheme like TDMA helps immensely by proving this fulfillment, which reduces development costs and time to market. Since most of the aircraft applications need a common time base, TDMA has the advantage that entities have to be synchronized to use it. This, at least, provides a relative time base.

Together with localization, five major prerequisites for WSNs for flight applications are identified: energy efficiency, transmission reliability, scalability, synchronization, and localization. Out of these five requirements, four are investigated in this thesis: the need for a common time base, the need to reduce power consumption, the need to increase transmission reliability, and the need to support several hundreds of nodes within one wireless cell.

2.2 Aircraft Radio Channel

All these applications have in common that they are using the same aircraft radio channel, which will be investigated in this chapter. The design of a WSN strongly depends on the physical characteristic of the environment in which it will be deployed. A WSN deployed in an industrial plant will face a different characteristic than one deployed in an outdoor area or in an aircraft.

To investigate the radio characteristics inside an aircraft several channel measurements were performed [9–12]. According to D’Errico et al. [12] two major factors influence the aircraft environment: first, simplified, an aircraft is a tunnel-like environment where the reflections of the walls yield an energy contribution which increases the channel gain. Following from this, the path-loss exponent γ in the log-distance path loss model (Equation (2.1))

$$PL(d)[\text{dB}] = PL(d_0)[\text{dB}] + 10\gamma \log_{10} \left(\frac{d}{d_0} \right) \quad (2.1)$$

would be smaller than 2 (the value for the free-space case). The second factor, however, is counteracting this. The seats and the partition between different areas of the aircraft act as electromagnetic absorbers; the identified path-loss exponents were always between 2 and 3.

Additionally, they concluded that a Rayleigh distribution approximates the fading component better than a Rice distribution because of the dense multipath environment; the line-of-sight (LoS) component is not predominant. Hence, they proposed a lognormal channel model with an Rayleigh-like fading coefficient X_σ :

$$\text{PL}(d)[\text{dB}] = \text{PL}(d_0)[\text{dB}] + 10\gamma \log_{10} \left(\frac{d}{d_0} \right) + X_\sigma[\text{dB}] \quad (2.2)$$

The channel model established by D’Errico et al. [12] is compared to different channel models from literature [9–11] and the free-space path loss model in Figure 2.5.

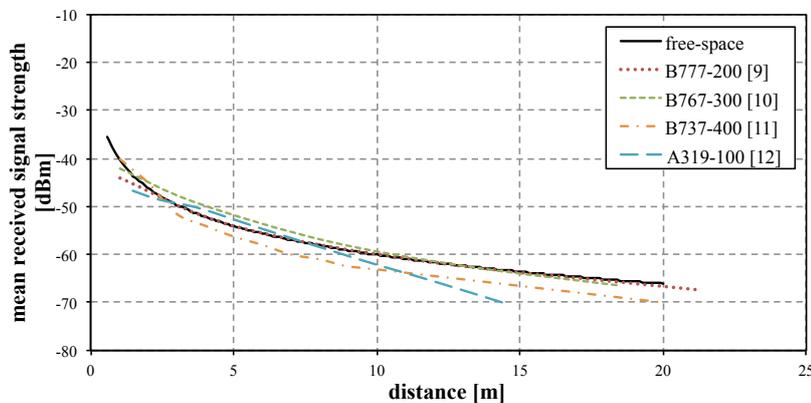


Figure 2.5: Aircraft channel measurements from literature [9–12] together with the free-space path loss model when transmitting with 0 dBm.

The different models predict a similar behavior. At a distance of around 10 m a path loss of around 60 dB appears. According to measurements [7], around -80 dBm is a typical power level where the packet error rate (PER) of an IEEE802.15.4 radio begins to increase significantly. To include a safety-margin in the system design, a power level of at least -60 dBm is desired. The length of a typical aircraft, for instance an Airbus A330, is around 60 m, hence, it becomes obvious that the whole aircraft cannot be served by only one access point (AP), when only single-hop communication is desired. To guarantee coverage of the whole aircraft, several wireless cells have to be deployed alongside the aircraft. In Figure 2.6 such a wireless cell distribution is depicted.

In the case of an Airbus A330, the Aircraft has to be compartmentalized into three wireless cells with a radius of about 10 m to cover the whole aircraft.

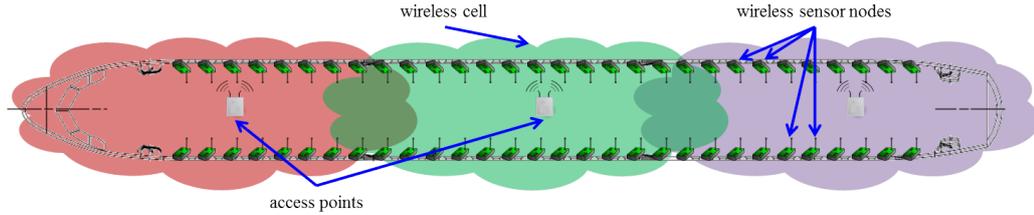


Figure 2.6: An aircraft covered by three cells using different channels.

To decrease the signal-to-interference ratio (SIR) between neighboring cells, each cell has to use a different channel, which can be derived from Figure 2.7.

In Figure 2.7(a) the path loss for two APs, AP 0 and AP 1, is shown. These APs are separated by a distance of $d = 50$ m. If both APs are using the same channel, the signal from AP 1 will interfere with the signal from AP 0, which is indicated by the $SIR(d)$ graph. At a distance of 25 m the signal strength of both APs is equal, hence, a receiving node at this distance from AP 0 would experience a $SIR(d = 25 \text{ m}) = 0$ dB. In Figure 2.7(b) such SIR graphs are shown for different separations between two adjacent APs.

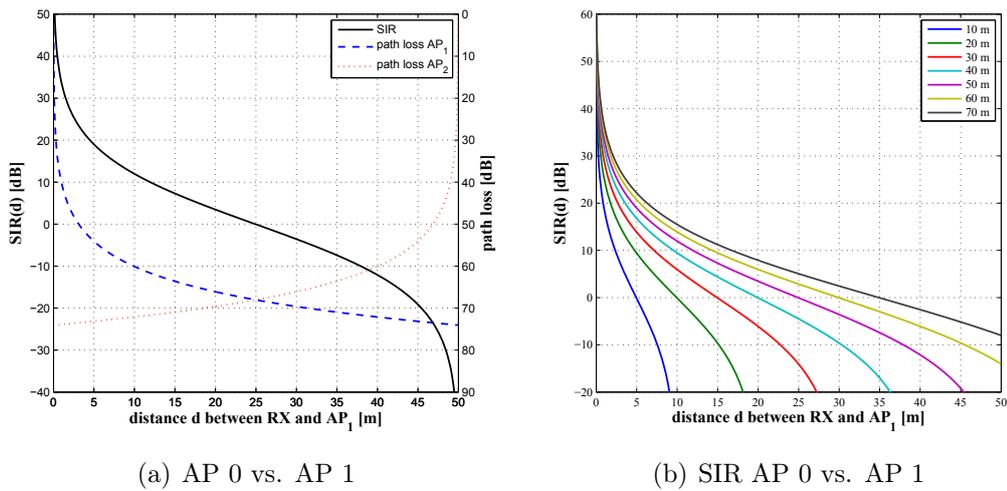


Figure 2.7: Aircraft channel re-use: (a): path loss off AP 0 and AP 1 separated by 50 m together with the resulting SIR between AP 0 and a receiver, and (b): SIR values between AP 0 and a receiver for several different distances between AP 0 and AP 1.

Petrova et al. [13] state that an IEEE802.15.4 radio can achieve a bit error rate (BER) of $1E-06$ with a signal-to-noise ratio (SNR) of roughly 14 dB. When examining a situation as depicted in Figure 2.6 where AP 0 and AP 2 are separated by 40 m, it can be concluded from Figure 2.7(b) that only a communication distance of 7 m is applicable for AP 0. At a greater distance

the SIR is greater than 14 dB, which would result in a higher BER than the intended $1E-06$. The wireless cells in Figure 2.6 were intended to cover a 10 m communication radius, hence, a frequency re-use is not possible in this case.

2.3 Duty Cycling

Another very important aspect for WSNs for flight applications is the need to have a very low power consumption. A common way to reduce the power consumption of a wireless node is to turn off all devices that are currently not used. When a transceiver is powered down it becomes impossible to initiate communication from the outside. Hence, to still support communication initiated from the outside the transceiver has to be kept active and has to listen to the channel. The receiver rarely receives actual data; for the majority of the time it will idly listen for a packet to arrive. To reduce the idle listening time and consequently the mean power consumption, the receiver will be turned off and on periodically. The percentage of the time the receiver is turned on is called duty cycle. With a decreasing duty cycle the power consumption in the receiver is decreased, but, at the same time, the mean communication delay is increased.

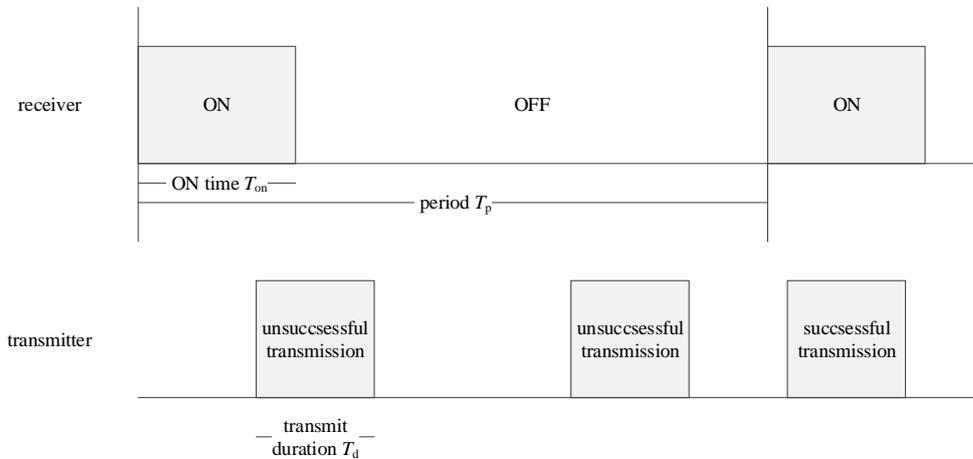


Figure 2.8: A 25 % duty-cycled receiver, together with a transmitter that is sending three packets until a successful transmission occurs.

In Figure 2.8 the periodic behavior of a duty-cycled receiver is depicted. The duty cycle follows $\frac{T_{on}}{T_p} \cdot 100\%$, where T_{on} is the time the transceiver is turned on and T_p is the duration of the period. In this case the receiver has a duty cycle of 25 %. The receiver is only able to receive a packet correctly if the transmit duration $T_d \leq T_{on}$ and the packet is transmitted during the ON time of the receiver, which means that only the third packet in Figure 2.8 can

be received correctly. It is only possible to communicate with a duty-cycled node in certain time frames; the period P equals the maximum additional communication delay before this frame arrives. Thus, there is a trade-off between power consumption and communication delay.

Additionally, when duty cycling, an overall more complex communication protocol is needed. To keep the communication delay as low as possible, the transmitting node either has to know beforehand at which time its communication partner is listening or it has to transmit the packet several times until it is certain that a listening slot of the receiving node is met. In the first case, the synchronous schedule, the packet has to be sent only once, but the clocks of all nodes have to be kept synchronous. In order to do this, additional communication has to be done which complicates the communication protocol and can result in higher power consumption for the transmitting node. In the second case, the asynchronous schedule, clocks do not have to be kept synchronous. In order to keep a low communication delay the packets have to be transmitted several times, which also increases the power consumption for the transmitting node. Hence, duty cycling clearly decreases the receiver power consumption but at the expense of not only an increased communication delay but maybe an increased power consumption for the transmitting node as well.

Following formula can be used to estimate the mean power consumption for a duty-cycled receiver.

$$P_{\text{mean}} = \frac{P_{\text{switch}} \cdot T_{\text{switch}} + P_{\text{on}} \cdot T_{\text{on}} + P_{\text{off}} \cdot T_{\text{off}}}{T_{\text{p}}} \quad (2.3)$$

When calculating the power consumption, the switching time T_{switch} from the OFF state to the ON state is not negligible. According to Equation (2.3) the mean power consumption for two commercially available state-of-the-art transceivers is calculated for varying duty cycles. The first transceiver considered is the Atmel AT86RF212 [14], which implements IEEE802.15.4 for the 700/800/900 MHz band. The second transceiver considered is the Nordic Semiconductor nRF51822 [15], which implements Bluetooth LE (2.4 GHz band). The second device is a System-on-Chip, which comprises a transceiver and a microcontroller but only the transceiver figures are taken into account. Both devices implement very different physical layers but both are optimized for low-power communication. For IEEE802.15.4 the 20 kb/s and 40 kb/s binary phased shift keying (BPSK) modulation as well as the 250 kb/s offset quadrature phased shift keying (O-QPSK) modulation is considered. For the nRF51822 the 1 Mb/s Bluetooth LE as well as the proprietary 250 kb/s and 2 Mb/s Gaussian frequency shift keying (GFSK) modulation is considered.

In Table 2.2 the characteristics for the different operation modes are shown. Regardless of the underlying physical layer, the sensitivity decreases with an increasing data rate while maintaining roughly the same power consumption. Within the duration T_{on} it has to be possible to transmit at least a packet

header with address information. Because of the different protocols this leads to a different quantity of necessary bits; 144 bits for IEEE802.15.4 and 80 bits for Bluetooth LE. Hence, the packet length and therefore the minimum ON time T_{on} is different for each protocol.

Table 2.2: Characteristics for commercially available transceivers

	data rate	sensitivity	P_{on}	P_{off}
AT86RF212	20 kb/s	-110 dBm	30.36 mW	0.66 μ W
	40 kb/s	-108 dBm	30.36 mW	0.66 μ W
	250 kb/s	-101 dBm	30.36 mW	0.66 μ W
nRF51822	250 kb/s	-96 dBm	41.58 mW	1.98 μ W
	1 Mb/s	-91 dBm	42.90 mW	1.98 μ W
	2 Mb/s	-85 dBm	44.22 mW	1.98 μ W

Using these characteristics from Table 2.2, in Figure 2.9 the mean power consumption of a duty-cycled IEEE802.15.4 radio [14] and of a Bluetooth LE radio [15] are depicted. With an increasing duty cycle period—while keeping the duration T_{on} constant—the mean power consumption for both radios decreases. Using a higher data rate decreases the sensitivity of the devices but does not change the power consumption. Hence, a shorter necessary ON time reduces the mean power consumption for a fixed duty-cycle period. Until a certain limit—a period of about 10s for the Bluetooth LE radio and around 100s for the IEEE802.15.4 radio—the mean power consumption curves decrease linearly in the log-log scale. After the limit, the basic load of the devices becomes predominant, the mean power consumption values approach the device-specific horizontal asymptote.

A similar predominance effect can be observed when increasing the data rate. The necessary ON time will be reduced accordingly, but the switching time T_{switch} does not change. Hence, the switching time will become the predominant factor when a certain data rate, respectively necessary ON time is reached. Further increasing the data rate will not significantly decrease the mean power consumption.

Hence, duty cycling a radio decreases its mean power consumption but only until a data rate dependent threshold is reached. Then the mean power consumption approaches an horizontal asymptote.

2.4 Wake-up Receiver Technologies

Another possibility to reduce power consumption is the use of specially designed hardware devices, which will be presented in this chapter. But only the protocols presented in Chapter 3 and 4 will use this hardware device.

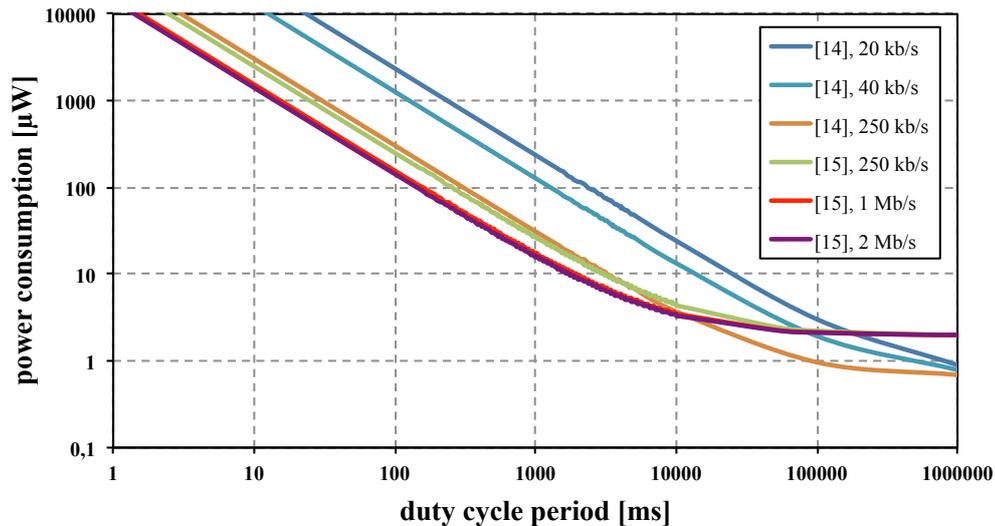


Figure 2.9: Mean power consumption of duty.cycled AT86RF212 [14] and nRF51822 [15] with different data rate settings.

Industrial process automation and infrastructure monitoring require dependable wireless sensor networks. These application areas demand not only energy efficiency but also reliable and timely data communication. Additionally, some application areas such as industrial automation or infrastructure monitoring impose strict latency bounds.

Duty cycling is the common way to reduce energy consumption but can lead to increased delay and a less reliable communication, which contradicts the requirements dictated by mission-critical applications. To overcome these limitations there are several efforts to develop low-power receivers together with appropriate protocols. The idea is to have receivers with such a low power consumption that it becomes feasible to keep them turned on all the time. Depending on the target application such a receiver can fulfill a variety of functions; among others, it can be used as a “classical” receiver or as a WUR. As the name indicates, a WUR is used to wake up an otherwise powered-down node to initiate appropriate tasks. Thus, a node will be woken only on demand and not periodically, which can decrease communication delay while keeping a very low power consumption.

2.4.1 Application Requirements

Depending on the application area, different requirements are demanded from low-power receivers. The requirements are mostly influenced by the communication range, the channel characteristics, the message flow, and the available energy. All these requirements have a direct impact on the key performance figures for low-power receivers. Communication range and channel

characteristics directly influence the necessary receiver **sensitivity**, and the message flow influences the necessary **data rate**. To improve either one of them, more energy has to be spent. Hence, the available energy is the limiting factor for all receiver designs; typically a very low **power consumption** is essential.

WURs can be used in several domains not only in the aeronautical one. But not in all circumstances will WURs have advantages over duty cycling a radio. Hence, to provide an overview where WURs can be used meaningfully and what their advantage over duty cycling a radio is, the requirements for four representative scenarios will be inspected: the aeronautical case, the wireless body area network, smart metering, and industrial applications.

Aeronautical Use Case

One challenge for wireless sensor networks in the aeronautical case is the high density of wireless sensor nodes [7]. The quantity of sensor nodes within an aircraft can easily reach 1000 or more. Out of the aircraft applications presented in Section 2.1 only Wireless Switch and Ambient Light Control have an event-triggered data flow. These event-triggered applications will benefit the most by using WURs, because of the non-predictability when the events occur. WURs do not have a periodic behavior as duty-cycled radios have. For these two applications, currently, duty-cycled transceivers are used to wake up the nodes at pre-determined points in time and then to start the protocol. For many applications, such a wake-up is only necessary once a day and therefore a very long duty cycle is used and a very low energy consumption is possible. Due to a very low duty cycle the introduced latency of the system is very high (> 10 s) while the responsiveness of the system has to be in the order of 200 ms, such a low duty cycle is not feasible.

To make the system autonomous, the power supply for the nodes is provided by an energy harvester [16] in combination with a power management board [17]. Such an energy harvester can provide up to 23 J per flight. With two flights a day and an efficiency of 54 % of the power management board [17], only 288 μ W are available on average. When the nodes are woken up, they have to start a measurement and transmit these values to a common sink. Measuring can consume a lot of energy, therefore, the available average power is further reduced to around 100 μ W.

Hence, the requirements for a low-power receiver in the aeronautical case are an average power consumption below 100 μ W, a sensitivity better than -80 dBm, and a delay below 200 ms.

Wireless Body Area Network

A wireless body area network (WBAN) consists of devices that are all located in close vicinity of the human body. Typical applications for WBANs are in

health care or in the multimedia sector. Latré et al. [18] give an overview of WBANs in general and the requirements resulting from their applications. As stated there, a WBAN will consist of around 20 – 50 sensor nodes, which are used to monitor bodily functions like heartbeat or body temperature, as well as to connect multimedia devices like microphones or head-mounted displays. Those devices can have very different needs for data rates in the region of only a few kb/s to several hundreds of kb/s; in sum the demands can reach a few Mb/s.

The advantage of low-power receivers in WBANs can be twofold: either a low-power receiver is used for low data rate applications and will replace a “normal” receiver, or it is used as a wake-up receiver to initiate a possible high data rate communication. Either way, the power consumption has to be very limited. For example, some health care devices are implemented inside the human body, like insulin reservoirs or pacemakers. It is clear that with such devices a regular battery change is no option. Additionally, because the devices are implemented inside the body, there is a strict size constraint for the nodes, e.g., it is not possible to use large batteries. Therefore, the power consumption has to be as low as possible.

The short communication distance of about 1 m that the devices have to cover can help to keep the power consumption low. Van Roy et al. [19] conducted a large measurement campaign to model the behavior of WBAN channels. A stochastic channel model was derived from it where each link is considered separately. The highest path loss was found around 50 dB, with all radio links on the front side of the body. With a path loss in the region of 40 – 50 dB the sensitivity requirements for low-power receivers are lessened; a sensitivity of –40 dBm would be enough to receive a signal transmitted with 0 dBm. Without the need to have a high sensitivity even more energy-efficient designs can be implemented.

Smart Metering

Another field of application for low-power receivers is in smart metering and home automation. In both cases, radio communication takes place inside a building and the requirements for both are similar. The maximum communication distance is typically about 15 m, but it has to be differentiated if the communication takes place on the same floor or across multiple floors. At a distance of 15 m for both cases the path loss is in the same order of magnitude but at shorter distances the multiple-floor scenario has a larger path loss than the same-floor scenario. Chrysikos et al. [20] characterize a channel for both cases. A path loss of around 100 dB was measured at a distance of 15 m. At 10 m a path loss of 100 dB was measured for the multiple-floor scenario and a path loss of 80 dB for the same-floor scenario. A similar behavior is described by Tsuchiya et al. [21] for the same-floor scenario.

Therefore, a low-power receiver for these applications has to have a high

sensitivity. With a maximum allowed effective radiated power (ERP) of 25 dBm in the European 868 MHz band, a receiver should have at least a sensitivity of -75 dBm to be able to receive packets at a distance of 15 m. But commercially available transceivers normally have an output power around 0 dBm. To generate an output power of 25 dBm typically a discrete power amplifier has to be used. Assuming a high efficiency of 50%, additional 28 dBm ($= 631$ mW) has to be spent to reach these high output power levels. Since a very high output power is rarely possible for low-power devices, a higher sensitivity should be used.

The required data rates for home automation / smart metering applications are moderate but the acceptable delay for a communication might be more demanding. If human interaction is involved, a latency larger than 0.5 s might not be tolerable. The allowed maximum energy consumption for the smart metering / home automation area depends heavily on the specific application. For some devices, like heat or electric meters, the battery has to provide enough energy for around 10 years, and because of that, large and expensive batteries have to be provided. Reducing the energy consumption of the communication system can reduce the overall costs immensely.

Industrial Applications

Control loops have the strictest requirements in the industrial applications field. Not only are they loss-intolerant but also delay-intolerant. For control loops, a maximum allowed delay in the order of 25 ms is common.

The wireless channel for industrial applications is similar to the one for the smart metering / home automation applications. The big difference is the target communication distance of 100 m. Tanghe et al. [22] characterize the industrial indoor channel. There, three types of topographies are distinguished: LoS, obstructed line-of-sight (OBS) with light surrounding clutter, and OBS with heavy surrounding clutter. For these three topologies channel measurements were done in three different bands, at 900 MHz, at 2.4 GHz, and at 5.2 GHz. From these measurements nine different channel characteristics were derived, three for each frequency band and three for each topology, respectively.

Tanghe et al. [22] argued that the log-distance path loss model is applicable for each characteristic and the corresponding path-loss exponents γ and the path loss values $PL(d_0)$ at a reference distance $d_0 = 15$ m were derived. When comparing the path loss at a distance of 100 m for the three topologies at 900 MHz, it was verified that the path loss behaved as assumed. The smallest path loss was observed in the LoS case with 98 dB, followed by the OBS with light surrounding clutter with 100 dB. The biggest path loss was observed in the case of OBS with heavy surrounding clutter with 110 dB. For the 2.4 GHz band the behavior was similar with 100 dB, 101 dB, and 112 dB, respectively. Assuming again a maximum transmit power of 25 dBm,

a minimum sensitivity of -85 dBm is necessary.

In industrial applications bigger devices and therefore bigger batteries might be possible, but the total cost of the system has to be kept as small as possible. A receiver that fulfills the sensitivity and delay requirements while fulfilling the strict energy consumption demands can help to lower the overall system and operation cost.

Summary

Depending on the application scenarios, requirements for low-power receivers differ. For some scenarios a very high sensitivity is necessary whereas in other scenarios a high sensitivity is not essential. As will be shown in Section 2.4.2, power consumption, sensitivity, and data rate are competing features and can be traded against each other; depending on the scenario these values have to be carefully chosen.

In Table 2.3 estimates for the system requirements for low-power receivers are given. These values are reference points for the system development and shall help to categorize low-power receivers; they are not meant to be strict, fixed limits for the system requirements. As can be seen, the highest requirement for sensitivity is given in the industrial case, closely followed by the aeronautical and smart metering applications. Only WBAN applications have a moderate sensitivity requirement owing to the short communication distance. The delay and data rate requirements are closely coupled; a very low data rate leads to a high communication delay which can be challenging. In some receiver concepts a data rate of only several hundred bits per seconds is possible; for instance, if 100 bytes have to be transmitted at a data rate of 1 kb/s this would lead to a minimum delay of 800 ms. Hence, the data rate of the low-power receivers have a direct impact on the possible minimum communication delay. The industrial case has the most stringent delay requirements. Regarding the data rate, WBAN applications have the highest demand because of possible multimedia applications and also the strictest requirements for power consumption.

Table 2.3: Requirements for low-power receivers

	distance	sensitivity	data rate	delay	power
Aeronautical case	10 m	-80 dBm	250 kb/s	200 ms	100 μ W
WBAN	1 m	-40 dBm	1000 kb/s	1000 ms	10 μ W
Smart metering	15 m	-75 dBm	10 kb/s	500 ms	100 μ W
Industrial	100 m	-85 dBm	250 kb/s	25 ms	200 μ W

Generally, industrial applications have the highest sensitivity and delay demands while still requiring a very low power consumption. Aeronautical applications have slightly lesser sensitivity demands but at the same time

stricter power consumption demands. Smart metering applications have lessened overall system requirements. Lastly, for WBAN applications a very low power consumption is clearly the most important factor, which can be met by providing only a moderate sensitivity.

2.4.2 Technology

Several concepts for low-power receivers are available [23–67], differing in complexity and performance.

The technologies used for low-power receivers vary over a wide range; from simple passive energy detection to envelope detection with several correlation stages. Each technology has its own advantages and disadvantages. Depending on the requirements for the low-power receivers, the best technology has to be selected carefully. Some are better for high data rates, others for high sensitivity or very low power consumption.

Modulation

The first distinctive feature is the chosen modulation format. Most of the concepts use either amplitude modulation in its binary form on-off keying (OOK), or they use frequency shift keying (FSK). Only four of the presented concepts differ here. The concept presented by Le-Huy et al. [59] uses binary pulse width modulation (PWM), while the concept by Marinkovic et al. [37] uses a multi-stage approach beginning with an OOK modulated signal followed by a PWM symbol. The other concepts use phased shift keying (PSK); Yan et al. [66] use the binary form BPSK and Lee et al. [60] use the quaternary form quadrature phased shift keying (QPSK), which could also be implemented as two parallel BPSK demodulators.

The advantage of FSK (and also PSK) over OOK is that it is less susceptible to noise and fading. Additionally, because of the constant power level of the carrier, the amplifier design is simpler with FSK and there is no need for an adaptable threshold to decide which symbol was received. OOK has the advantage of overall implementation simplicity which can be transferred into energy efficiency. For instance, just a simple envelope detector can be used to detect if the carrier frequency is turned on. Such a detector can easily be implemented by using a diode and a resistor-capacitor oscillator circuit as a bandpass filter. All the concepts that have a power consumption below $10\ \mu\text{W}$ are using OOK. This can be seen in Figure 2.10.

Implementation

The second distinctive feature is the implementation of the concept. Several variants are used: super-regenerative receivers, superheterodyne receivers, injection-locked local oscillators (LO), envelope detectors with or without an

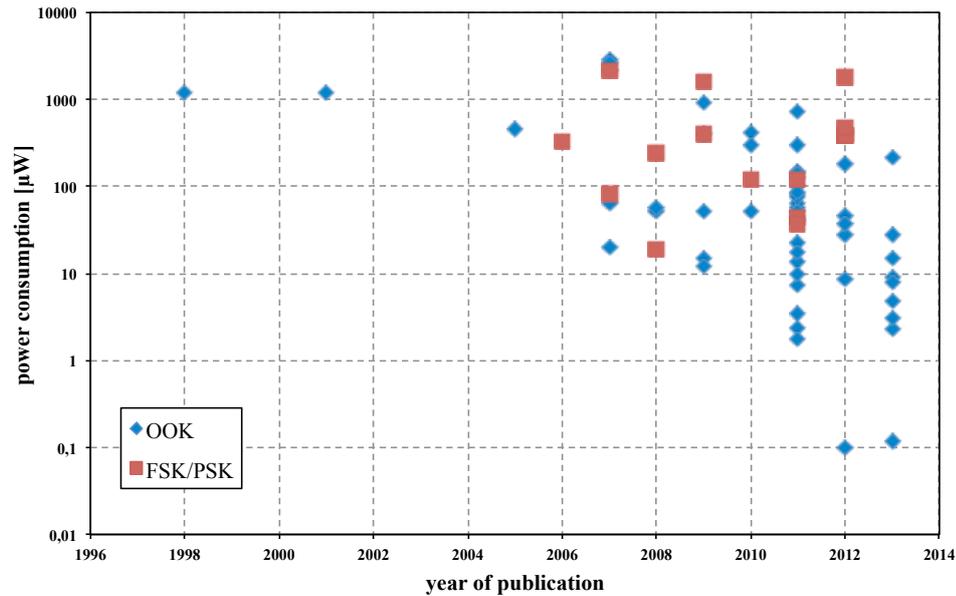


Figure 2.10: Power consumption of the low-power receivers using the OOK [23–56] and FSK/PSK [57–67] concepts.

intermediate frequency (IF), and in two cases envelope detectors followed by purely analog correlation stages.

The concepts in [23, 24, 29, 36, 41, 45, 52, 53, 56, 57] are using a super-regenerative receiver design. They all have a data rate above 100 kb/s with a sensitivity varying over a broad range. Two designs [31, 44] are using envelope detectors followed by purely analog non-coherent correlators and, thus, are able to reach very low power consumption values. Five concepts are using duty cycling [27, 34, 38, 39, 55], while some concepts [27, 38, 39, 55] are duty cycling even within a bit having an ON time in the range of only 100 ns. This reduction in sampling time decreases the power consumption to a level of several μW at a low data rate (< 1 kb/s) without increasing the latency due to duty cycling. The concepts in [54, 62, 66, 67] are using injection-locked oscillators for either creating a power efficient “virtual LO” or creating a power efficient FSK / PSK demodulator. The concepts in [27, 30, 40, 42] are using an uncertain IF architecture where the LO can vary in the range of 100 MHz from its desired frequency and therefore can be implemented more power-efficiently. In [34, 50, 61, 66] no IF is used, the symbols are demodulated directly which again can simplify the receiver design immensely. A very power-efficient design is possible but at the cost of a low sensitivity and the need of an external surface acoustic wave (SAW) filter.

Regardless of the particular specifications of the designs, the commonly important characteristics of all low-power receivers are **power consumption**, **data rate**, and **sensitivity**. None of these technologies can clearly be

identified to outperform the others in all of these characteristics. All parts of a receiver strongly influence its performance; the interaction between them accounts for the overall power consumption. Therefore, all the low-power receivers will be compared with each other regardless of the specific technology.

2.4.3 Comparison

Sensitivity vs. Power Consumption

First, the power consumption of the low-power receivers is related to their sensitivity. If possible, the sensitivity is measured at a BER of 10^{-3} ; due to correlation stages, for some concepts the sensitivity can only be measured at PER levels because their functional principle makes no information available at BER level. This has to be kept in mind when comparing the sensitivity levels. The supply voltages vary in the range of 0.4 V–3 V. Some concepts use multiple supply voltage levels, meaning the actual power consumption might be higher due to the necessity of voltage converters when implementing the concepts on a communication platform. The behavior of some concepts is adjustable, therefore, these receivers will have several connected data points in the figures. The data points of some concepts— [38,39] in Figure 2.11 and [33,62] in Figure 2.12—describe a vertical line, suggesting that one receiver setting outperforms all the others. However, this advantage is only valid for the examined parameters. As can be seen in the corresponding figure, there is no receiver setting which outperforms all the others. Nevertheless all data points are drawn to indicate that several receiver settings are feasible, even if they draw a vertical line. As can be seen in Figure 2.11, it is very hard to reach a power consumption smaller than $10\ \mu\text{W}$. Only nine concepts [31,37–39,43,44,49,54,55] are below that level, where [44,49] reach an outstanding power consumption around $100\ \text{nW}$. All those concepts can meet the moderate sensitivity demands for the WBAN applications.

If both a small power consumption and a high sensitivity shall be maintained the situation becomes even more challenging. As can be seen, only three concepts [32,39,60] consume less power than $100\ \mu\text{W}$ at a sensitivity better than $-80\ \text{dBm}$, and hence, fulfilling the requirements for the aeronautical applications. For industrial applications only two concepts [60,62] meet the requirements of a sensitivity better than $-85\ \text{dBm}$ while maintaining a power consumption below $200\ \mu\text{W}$.

Data Rate vs. Power Consumption

Second, the data rate of the receivers is related to their power consumption. Depending on the concepts, different coding techniques are used. Some designs [28,43,51] use Manchester coding to be able to recover the clock signal from the data. With Manchester coding the net data rate is only half of the

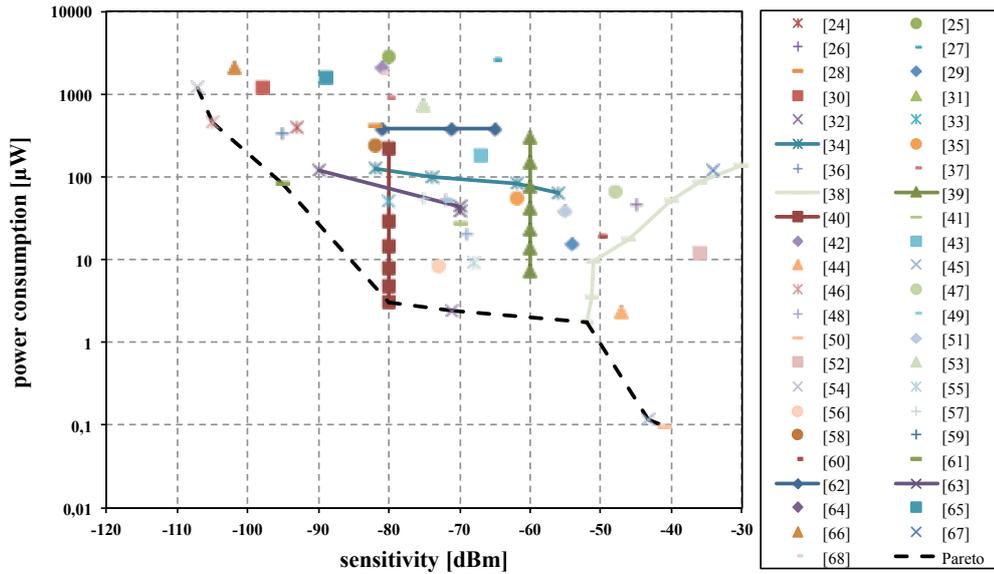


Figure 2.11: Sensitivity of the low-power receivers vs. their power consumption together with Pareto front.

gross data rate. Other designs [25, 31, 37–39, 43, 44, 46, 52, 59] use spreading techniques to increase the sensitivity or to distinguish between codes. Here, the net data rate is in the order of the spreading factor smaller than the gross data rate. In the figures, gross data rates before spreading, correlation, coding, etc. are shown.

As can be seen, it is hard to reach high data rates while keeping a low power consumption. Only ten designs [26, 30, 41, 42, 50, 52, 57, 63, 66, 67] reach a data rate of at least 1 Mb/s. From those only three [42, 50, 66] have a power consumption smaller than 200 μW .

Regarding the requirements in Table 2.3, it can be seen that for the aeronautical and the WBAN case no low-power receiver concept fulfills the power consumption, sensitivity, and data rate requirements. For the industrial case only [62] and for the smart metering case only [32, 56] fulfill all three requirements at the same time. The characteristics of these concepts would be sufficient to fulfill the sensitivity and data rate demands of those applications while reducing the power consumption dramatically.

For the aeronautical and for the WBAN scenarios currently no low-power receiver concept fulfills all the requirements to replace the main receiver of the system. However, a low-power receiver implemented as a wake-up receiver can have advantages over duty cycling a main transceiver: either the overall power consumption and the communication delay can be reduced or at least the communication delay can be reduced while maintaining the same power consumption a duty-cycled transceiver would generate. Section 2.4.6 addresses this aspect more thoroughly.

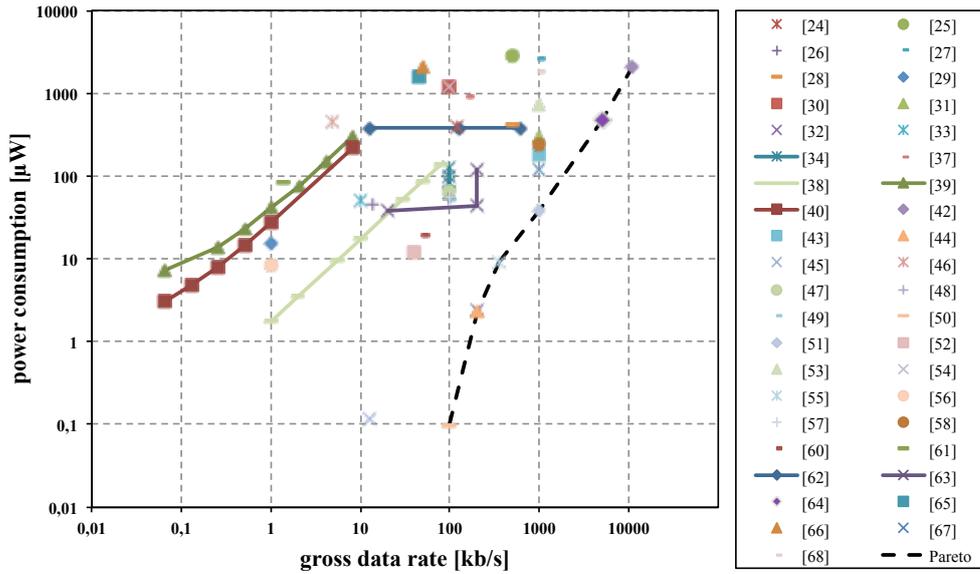


Figure 2.12: Gross data rate of the low-power receivers vs. their power consumption together with Pareto front.

2.4.4 Implementation Examples

Two designs are presented in more detail: the WUR design by Hambeck [68] and the WUR design by Milosiu et al. [39]. These two designs were integrated in wireless sensor nodes and form the basis for the evaluations in Chapter 3 and Chapter 4.

Wake-up Receiver from the CHOSeN project

The first wake-up receiver design presented is developed by Hambeck [68] in the context of the European Union (EU) project CHOSeN [69]. The application-specific integrated circuit (ASIC) uses a direct down conversion architecture and includes an envelope detector, a low-noise amplifier (LNA), a programmable-gain amplifier (PGA), and a mixed-signal correlation unit.

In Figure 2.13 the block diagram of this wake-up receiver is depicted. It is necessary to use an external SAW filter for out-of-band interference compression as well as a high quality matching network for voltage transformation. By carefully designed building blocks the 868 MHz WUR only consumes $2.4 \mu\text{W}$ at -71 dBm sensitivity. This is achieved by baseband correlation over 7 ms using 64 bit long OOK patterns.

This analog correlation unit is depicted in Figure 2.14. As can be seen, it consists of 128 parallel sample gates and one-bit correlation stages. The known pattern against which the incoming signal will be compared is stored in a circular pattern shift register. The incoming signal will be sampled and each sample charge will, depending on the actually configured bit per correlation

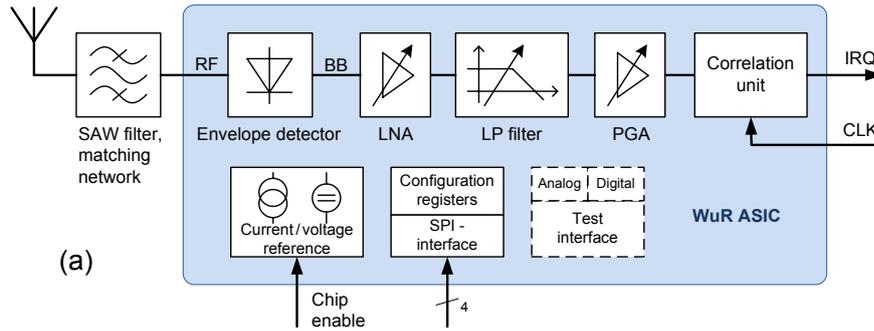


Figure 2.13: From [31]: block diagram of the wake-up receiver from the EU project CHOSeN.

stage i , either load the capacitor $C_{p,i}$ for positive bits or the capacitor $C_{n,i}$ for negative bits. After the duration of half a bit, each correlation stage will use the next bit of the pattern. This is done parallelly 128 times to, firstly, oversample the 64 bit pattern and, secondly, to guarantee that at least one of the stages will be synchronized to the incoming pattern. The correlation result is then evaluated once per pattern period by the comparators. If one or more of the comparators generate a high output, then a wake-up interrupt will be raised.

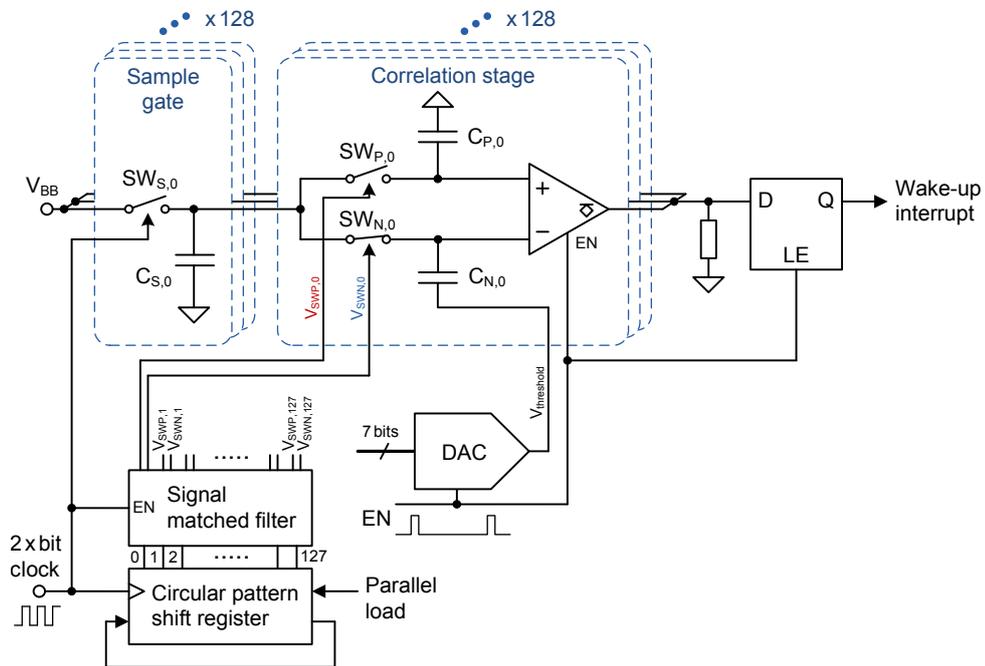


Figure 2.14: From [68]: schematic of the correlation unit of the wake-up receiver from the EU project CHOSeN.

It becomes possible to distinguish between several different patterns when using a set of orthogonal patterns. A kind of wake-up ID can be transmitted that is normally used for addressing the node that shall be woken up. Each ID is represented by one specific pattern out of the orthogonal set of patterns. Hence, the length of this ID equals the cardinality of the set. Since the wake-up receiver design by Hambeck [68] is only capable of comparing an incoming signal against one specific predefined pattern, only those WURs tuned to this incoming pattern will be woken up; WURs tuned to another pattern will not recognize the pattern at all. Thus, no continuous data reception is possible; only one specific pattern can be recognized. When another ID shall be received the WUR has to be tuned to listen for another pattern.

The length of this wake-up ID is limited due to the small set of available 64 bit patterns with good correlation properties. Promising code-sets with low cross-correlation are Gold codes or Kasami codes [70, 71]. Gold codes have a code family size of 65 and Kasami codes have a family size of 520. It can be shown that the largest cross-correlation value R_{\max} of Kasami codes is lower than the largest of Gold codes but the cross-correlation of Gold codes is lower on average than the one of Kasami codes. Zeng et al. [72] summarizes some code set characteristics, depicted in Table 2.4.

Table 2.4: Some Code Sets of length $(2^n - 1)$ for even n

Family	Family size	R_{\max}
Gold codes	$2^n + 1$	$1 + 2^{(n+1)/2}$
Kasami codes—large set	$2^{3n/2} + 2^{n/2}$	$1 + 2^{n/2+1}$

Wake-up Receiver from the AETERNITAS project

The second wake-up receiver design presented is developed by Milsoiu et al. [39] in the context of the Bundesministerium für Bildung und Forschung (BMBF) project AETERNITAS [73]. It comprises an analog superheterodyne front-end and two digital 31 bit correlation units. It uses a free-running LO for the down-conversion and does not need a phase-locked loop (PLL). Figure 2.15 depicts the block diagram of this wake-up receiver design.

The 868 MHz WUR achieves -83 dBm sensitivity at a scalable mean power consumption of down to $3 \mu\text{W}$. This very low power consumption is achieved by duty cycling the WUR during the reception of each singular data bit. This instance can be seen in Figure 2.16. In the top the received OOK modulated data bits are shown, in the bottom the asynchronous oversampling is shown. Each data bit is four times oversampled and between sample points the analog blocks of the ASIC are turned off. Usually, a receiver would be turned on during the reception of a whole data packet, respectively data bit. The novel approach to duty cycle between sample points reduces the mean power

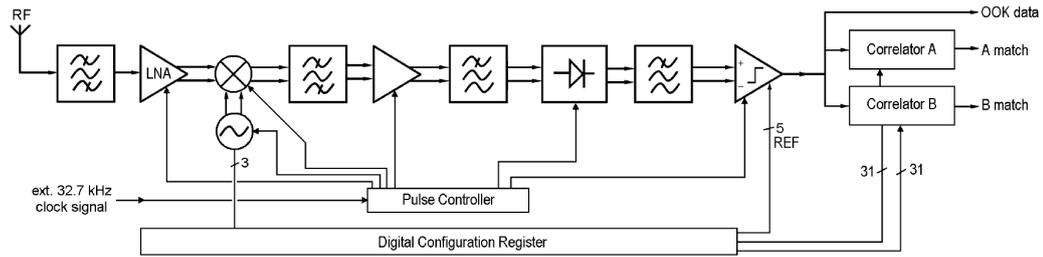


Figure 2.15: From [39]: block diagram of the wake-up receiver from the BMBF project AETERNITAS [73].

consumption immensely. To achieve this, the settling time of the analog devices has to be very short—around 100 ns instead of several hundred μs (370 μs for the AT86RF212). Otherwise, it would be impossible to use such a low duty cycle within one data bit.

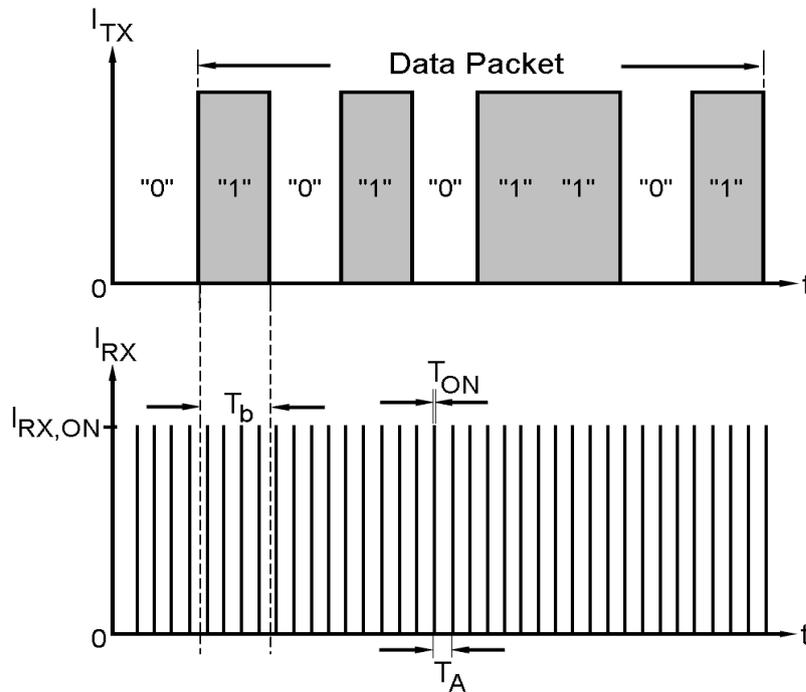


Figure 2.16: From [39]: asynchronous oversampling of OOK modulated data bits.

The settling time T_s and the sampling duration T_{ON} are constant, but by varying the sampling period T_A the duty cycle can be changed and with it the actual data rate. When the data rate is doubled the bit duration is cut in half, therefore, the duty cycle and with it the power consumption is nearly doubled as well. The sensitivity of this WUR is constant because it depends only on the duration of T_{ON} which is kept constant as well. The data rate

can be scaled in eight steps between 64 b/s to 8192 b/s, which scales the power consumption between 3 μ W and 217 μ W, respectively. This enables the possibility to adapt the WUR system to specific application requirements.

The WUR design comprises two digital correlation units which continuously compare the received bit strings against two different wake-up patterns. If the correlation result of either Correlator A or B exceeds a pre-configured threshold the corresponding WUR interrupt signal will be created. This signal can either be used directly as a wake-up command, or both the signals can be interpreted as binary data bits itself. Then, the net data rate would be reduced by a factor of 31 but a continuous data reception becomes possible, the wake-up receiver can be used as a very low-power receiver.

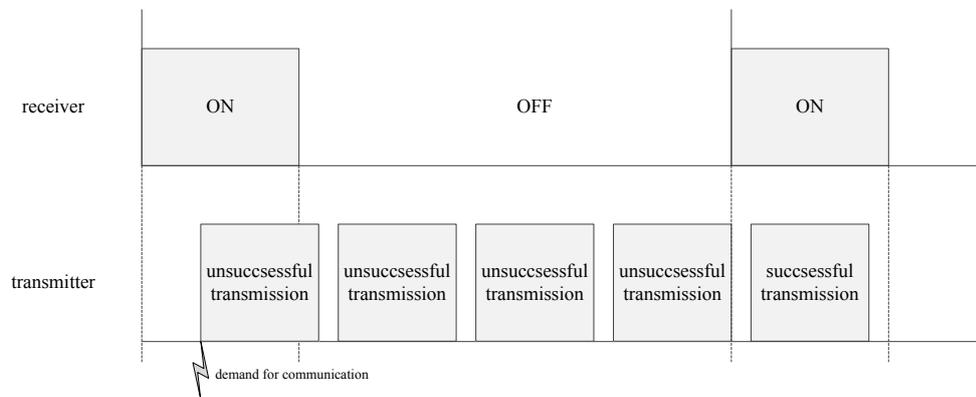
2.4.5 Application Areas for Wake-up Receivers

The most promising application for low-power receivers is to replace a “normal” receiver. If used for all communication purposes, the low-power design can yield the most energy savings. However, as outlined in Section 2.4.3, low-power receivers do not meet all requirements for all applications. Another operational area is to use a low-power receiver as a wake-up receiver. The low-power receiver will not be used for communication itself but only for initiating the communication via another receiver. Since the main transceiver does not have to be periodically duty-cycled, it can be kept turned off as long as no communication is required; it will be turned on only on demand. This behavior can be seen in Figure 2.17.

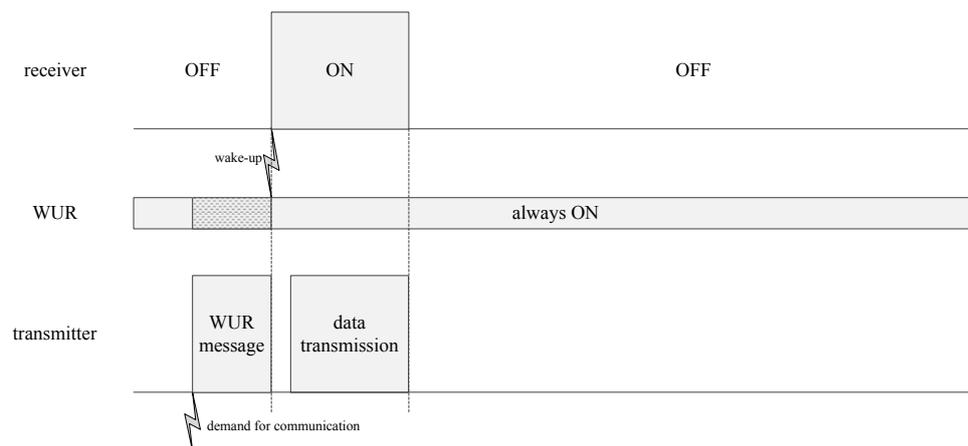
In the first case, the unsynchronized communication with a duty-cycled receiver is depicted in its simplest form. Once a communication demand arises, the transmitter wants to start communicating and repeatedly sends the data packet. This repetition must continue for the length of the duty cycle period to assure the meeting with the ON phase of the duty-cycled receiver. In the second case, instead of sending the data packet directly, first a wake-up message is transmitted. This message indicates the receiving node that a data packet will follow and therefore it will turn on its receiver. In this case, using a wake-up receiver will not only decrease the idle listening time but also shorten the mean communication delay.

This is beneficial in a typical aeronautic application: on the ground, the status of simple switches has to be evaluated. To accomplish this, the monitored switch can be connected to a sensor board, which currently is using a duty-cycled transceiver. With this system, the status of one switch can be polled within one second while only consuming $< 100 \mu$ W on average. Using a wake-up receiver instead, the communication delay could be decreased by approximately one magnitude at the same power consumption. This increases the responsiveness of the system immensely; it is possible to inquire the status of several switches in the same time as only one with a duty-cycled transceiver.

Polling devices is not the best solution for a large number of devices. Gordon et al. [74] present a more sophisticated protocol where acknowledged wake-up messages are used to initiate, for example, a TDMA protocol. In doing so, the time to inquire the status of a large amount of sensors can be reduced roughly by half. Another approach will be presented in Chapter 4.



(a) duty-cycled receiver



(b) wake-up receiver

Figure 2.17: Different communication approaches: (a) when using a duty-cycled transceiver, and (b) when the communication is initiated with a WUR message.

2.4.6 Wake-up Receiver vs. duty-cycled Receiver

Both duty-cycled transceiver and wake-up receiver have advantages over each other in certain areas. In the following, the boundaries of these areas are outlined.

In Figure 2.18 the power consumption of commercially available transceivers is depicted together with wake-up receivers with a power consumption below $200\ \mu\text{W}$. The delay in the case of WURs is equal to the time it takes to send a wake-up message. For designs [25, 31, 37–39, 43, 44, 46, 59] that use spreading techniques with fixed spreading factors, the amount of bits to send is predefined. For other designs, we assume that a wake-up message consists of 32 bits with the delay being directly proportional to their data rates.

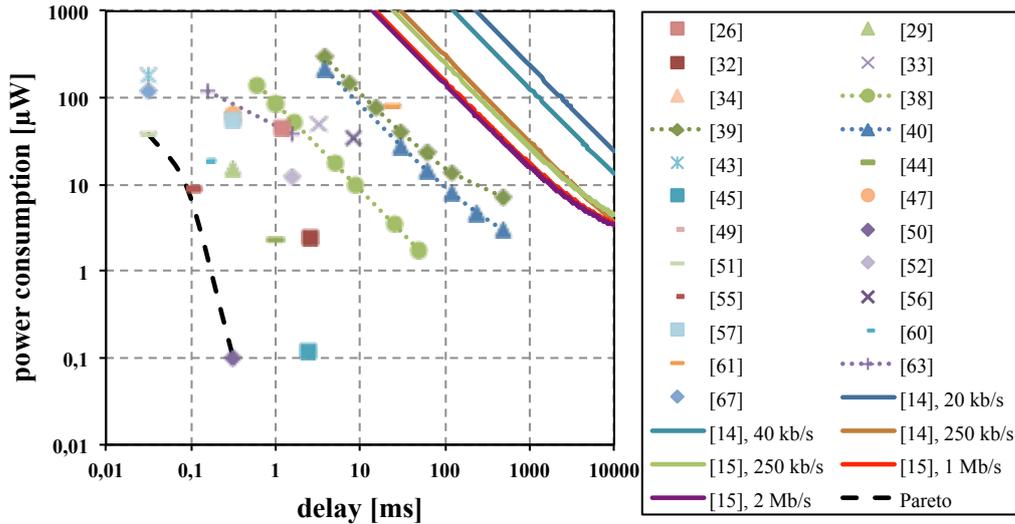


Figure 2.18: Wake-up delay of low-power receivers [25, 28, 31–33, 37–39, 42–44, 46, 48–51, 54–56, 59, 60, 62, 66] with a power consumption below $200\ \mu\text{W}$ together with duty-cycled transceivers Atmel AT86RF212 [14] and Nordic Semiconductor nRF51822 [15]; Pareto front included.

As can be seen, all WUR concepts perform better than the commercially available transceivers regarding delay and power consumption. All of the concepts fulfill the delay requirements given in Table 2.3. Considering the sensitivity requirements as well, only a subset remains. For the aeronautical case only three concepts [32, 39, 60] fulfill the requirements, for the industrial only one [60]. A duty-cycled nRF51822 with a data rate of one or two Mb/s also meets the industrial requirements, but only with a sensitivity worse than the concept in [60].

For an application with a delay restriction of 1 s, a duty-cycled transceiver is a sound option. The mean power consumption is in the order of $10\ \mu\text{W}$ to $100\ \mu\text{W}$ with a sensitivity of $-90\ \text{dBm}$ to $-110\ \text{dBm}$. This increase in sensitivity can improve the reliability of the communication system which outweighs the additional effort for the more complex communication protocol.

Conclusions

Development over the last 10 years clearly improved the performance of low-power receivers. The power consumption dropped from 1 mW to almost 100 nW, hence, a factor of 10,000. Still there are no designs available which fulfill the combined requirements for energy consumption, sensitivity, and data rate for the aeronautical and for the WBAN applications. If the performance improvement continues to proceed, concepts for all application areas should be available soon.

Currently there are few applications which depend on the usage of a low-power receiver / wake-up receiver. Many applications could clearly benefit from using a wake-up receiver; either because of a reduced power consumption or because of a smaller communication delay. Some applications will only be enabled when using a low-power receiver. With the high potential for reducing the energy consumption, protocols that are dependent on wake-up receivers are expected to arise.

Depending on specific application requirements, it is already possible to replace a “normal” receiver with a low-power receiver. If the low-power receiver characteristics do not yet meet the requirements for the applications, they can be used as wake-up receivers, which, by keeping similar power consumption values, drastically reduces the communication delay in comparison with duty-cycled systems.

3 Energy-Efficient Clock Synchronization

In this chapter an energy-efficient way to create a common time basis is presented: wake-up receiver messages are utilized to synchronize clocks energy-efficiently.

3.1 Introduction

Many applications and protocols in the WSN domain require a common time basis. Amongst others, this time basis is necessary to enable some MAC schemes, to determine the occurrence of special events, or to start measurements at the same time. But the clocks of the sensor nodes have to be synchronized as energy-efficiently as possible.

Various protocols are developed to keep clocks synchronized, but they either are too power-demanding or require a special MAC protocol. A simple, very energy-efficient, robust synchronization protocol is presented which bounds the synchronization error to a certain limit and which does not need a special MAC protocol.

3.2 Background

A typical wireless sensor node, as depicted in Figure 3.1, comprises a microcontroller, a transceiver, an antenna, a sensor, and a clock oscillator. Its microcontroller and its clock oscillator are used to generate the local time of a sensor node. The clock oscillator generates clock pulses which are used to increment a counter in the microcontroller. This counter value then represents the local time t of a sensor node. Hence, the quality of the local time strongly depends on the used clock.

In the following background information is given about clock oscillators, periodic clock synchronization, possible delay sources, and about synchronization protocols.

3.2.1 Clock Oscillators

A typical clock oscillator for WSNs consists of a crystal oscillator and of an associated oscillator circuit. The crystal oscillator is tuned to oscillate on one

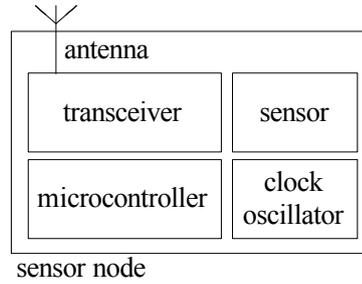


Figure 3.1: Wireless sensor node with its typical components.

specific frequency and provides a sinuous output, which is transformed by the oscillator circuit to a rectangular one. Additional control functionalities are usually also provided by the oscillator circuit. The characteristics of a clock—the nominal frequency f_0 , the frequency tolerance in parts per million (ppm), and the frequency temperature characteristic—determine the behavior of the local time. The nominal frequency is the frequency at which a clock oscillator is intended to run. This frequency will vary in the region of the frequency tolerance due to fabrication-induced imperfections. Additionally, the clock frequency depends on temperature. For instance, a typical 32 kHz tuning fork crystal oscillator [75] is designed to have a parabolic frequency-temperature characteristic with a maximum turning point at 25°C. At this temperature the crystal has its nominal frequency; with decreasing as well as with increasing temperature the actual frequency will decrease. In Figure 3.2, this characteristic is shown for the clock oscillator OV-7604-C7 [75].

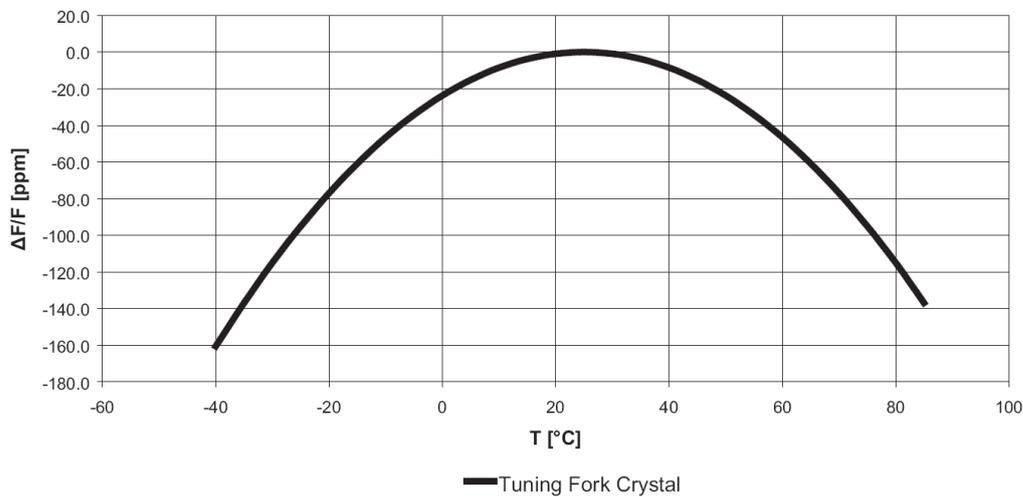


Figure 3.2: From [75]: temperature-dependent frequency characteristic of a tuning fork crystal.

The nominal frequency, the frequency tolerance, and the frequency temperature characteristic determine the instantaneous frequency $f_i(t)$ at which

the counter in the microcontroller will be increased. The local time of Node i can be described by $T_i(t)$. The derivative of T_i , the clock rate, describes the rate at which a clock is progressing. If Node i has a perfect clock ($f_i(t) = f_0$) the following holds: $T_i(t) = t + c$. Then, the clock rate $T'_i(t) = 1$. Hence, the local time increases as fast as the real time. In Figure 3.3, the local time $T_i(t)$ of two nodes is depicted: Node 1 with a 10 % slower clock and Node 2 with a 20 % faster clock.

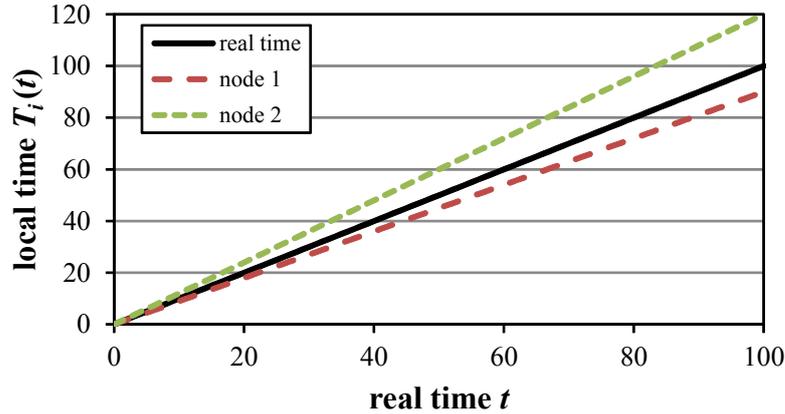


Figure 3.3: Local times for two nodes; Node 1 with a 10 % slower clock ($T'_1 < 1$) and Node 2 with a 20 % faster clock ($T'_2 > 1$).

As can be seen, the local time of both clocks drift apart from the real time. The difference between the local time to the real time at time t is called clock offset and follows $\Delta T_i(t) = T_i(t) - t$; the clock offset of Node i relative to Node j follows $\Delta T_{i,j}(t) = T_i(t) - T_j(t)$, with $i \neq j$. Accordingly, the difference of clock rates is called clock skew and follows $\Delta T'_{i,j}(t) = T'_i(t) - T'_j(t)$, with $i \neq j$. Clock skew is a measure for how fast the local time of two clocks will drift apart. For instance, Node 1 in Figure 3.3 has a clock skew of $\Delta T'_1 = -0.1$ and Node 2 of $\Delta T'_2 = 0.2$. The relative clock skew between Node 1 and 2 equals $\Delta T'_{1,2} = -0.3$, which means that the difference of the local times of Node 1 and 2 will increase by 0.3 seconds per second. The resolution of the local time is quantized. A tick of the counter in the microcontroller corresponds to a time quantile and will be increased with each clock pulse. Hence, the clock resolution equals the clock period $\frac{1}{f_0}$.

Several strategies are possible to align the local time: adjusting the counter in the microcontroller to match the real time, adjusting the crystal frequency to compensate for temperature-induced frequency offsets, or aligning the local time “off-line” where the time offset is stored but not adjusted. Temperature-induced frequency offsets of some clocks can be adjusted. A precise but costly method is to use an oven-controlled crystal oscillator (OCXO). The crystal of an OCXO is encased in a temperature-controlled chamber which keeps it at

constant temperature, hence, the temperature-induced frequency offset will be minimized. This method leads to very precise clocks but has the drawback of a high power consumption and of a rather big size. Thus, OCXO are seldom used in WSNs. Another approach is not to stabilize the temperature of the crystal but to adjust its frequency output directly. The output frequency of an voltage-controlled crystal oscillator (VCXO) can be adjusted by few ppm over a typical voltage range of 0 V to 3 V. If the frequency difference is only in this magnitude and can be obtained with enough precision, then a VCXO can be a valid option. But it is often hard to acquire a precise momentary frequency offset. Additionally, a more complex circuit design has to be done to integrate a VCXO in a wireless node. In addition to connecting the clock signal of the VCXO to the microcontroller, a controllable voltage source has to be connected to the VCXO as well. This controllable voltage source also increases the total power consumption of the sensor node. Hence, typically just the value of the counter in the microcontroller is adjusted, which keeps track of time.

3.2.2 Periodic Synchronization

To compensate for clock drift and to maintain a common time basis clocks have to be synchronized. Therefore, each node has to have an additional time reference $T_r(t)$ to which the local clocks can be synchronized. There are many possibilities to acquire a time reference. A Global Positioning System (GPS) receiver [76] can provide a very precise time reference but is costly and has a high energy consumption. Another possibility is to exchange information about the time reference over the wireless channel, which can be more energy efficient but also more imprecise.

The time reference $T_r(t)$ provides a reference time for certain points in time t_r ; the local time can be aligned to the real time at time t_r : $T_i(t_r) = T_r(t_r)$. Most often the frequency offset will not be compensated for. Instead, the local time will be adjusted periodically. In between two synchronization points the local time is allowed to drift away. The synchronization period δ_s together with the relative clock skew determine how large the clock offset can get. When assuming a constant clock frequency the following holds for the clock offset:

$$\max \Delta T_i(t) = \delta_s (T'_i - 1) \quad (3.1)$$

This can be seen in Figure 3.4 where the local time for one node is shown. Without synchronization the offset steadily increases with time, with synchronization the offset increases only until a new synchronization point is arrived at which it is minimized again.

The accuracy of the synchronization depends on the accuracy of the reference signal and is bounded by the local clock resolution. In the following,

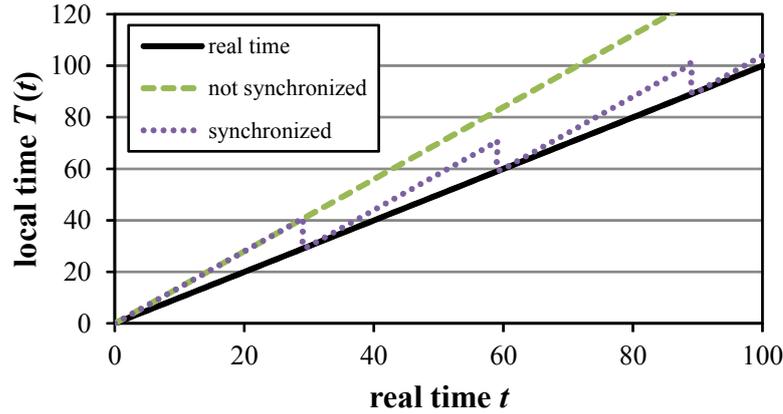


Figure 3.4: Periodically synchronized local time with a synchronization period of $\delta_s = 30$ s without synchronization errors.

we only consider time reference signals obtained by wireless communication.

3.2.3 Delay Sources

In order to calculate the offset between two clocks, the local time of both clocks at a certain common point in time has to be known. One possibility to gain this information is to transmit an actual clock reading $T_1(t_r)$ from Node 1 to Node 2. Node 2 then has to estimate the delay δ_{total} between the generation of the reading in Node 1 and the processing of the clock reading in Node 2. With this information Node 2 can synchronize its clock to the clock of Node 1: $T_2(t_r + \delta_{\text{total}})$ is set to $T_1(t_r) + \delta_{\text{total}}$, where $T_2(t_r + \delta_{\text{total}})$ is the time when Node 2 processes the clock reading of Node 1. The delay δ_{total} consists of the sum of several delay sources with different characteristics: the medium access delay δ_{mac} , the propagation delay δ_p , the transmission delay δ_t , and the operating system delay δ_o . In Figure 3.5 the components of delay δ_{total} are depicted.

The delay δ_{mac} is the time it takes Node 1 to access the wireless medium. Its jitter strongly depends on the used MAC scheme and on the network load. For example, when using CSMA it has to be established that a medium is currently not used before it is accessed. This process induces a random time variable which is network load dependent and hard to estimate. The delay δ_p is the time a radio wave has to travel from Node 1 to Node 2 and depends on the distance between the nodes. If the distance is known, this delay can easily be calculated. Additionally, for short distances (smaller 50 m) this delay is a magnitude smaller than the other delay sources and can be neglected. The transmission delay δ_t is the time it takes to transmit the whole data packet. Its mean value depends on the length of the data packet and on the used data rate. If both are known, the mean of the delay can be estimated with

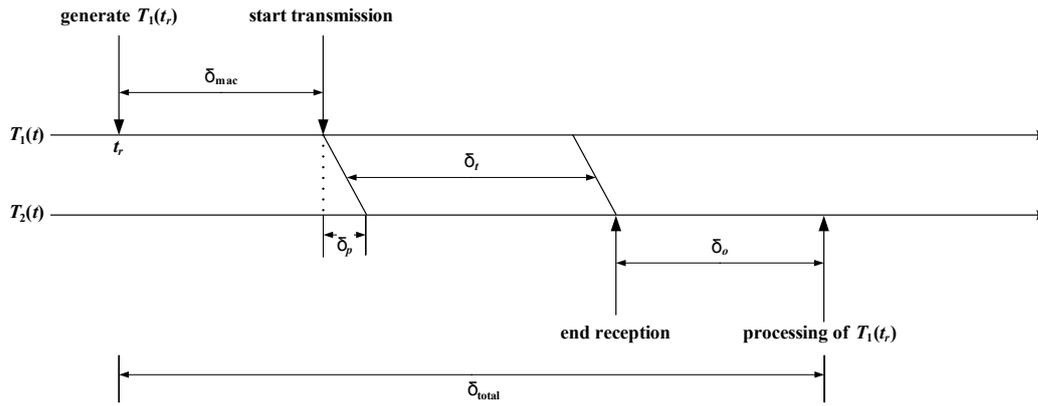


Figure 3.5: Composition of the delay δ_{total} .

high precision. But the variance of this delay depends on the used transceiver technology and is hard to estimate. Lastly, the operating delay δ_o is the time it takes the microcontroller in Node 2 to begin processing the data packet. This time strongly depends on how the microcontroller is programmed, for instance, if an operating system with task scheduler is used. Depending on the programming, it becomes very easy or very hard to estimate this delay.

3.2.4 Clock Synchronization Protocols

Several protocols [77–89] were developed to maintain a synchronized time in a network. Sivrikaya et al. [90] and Sundararaman et al. [91] give thorough overviews of them.

Out of the stated protocols, three [77, 80, 81] are now explained in more detail, because they present an important development of synchronization protocols for aircraft applications.

Mills [77] presents Network Time Protocol (NTP), which is commonly used in the Internet. NTP depends on the assumption that the delay δ_{total} is stable for a short period of time. A two-way message transmission is used to estimate the delay without further knowledge about the underlying processes. A node that wants to synchronize itself to a server transmits a synchronization request to it. The server answers with its actual clock reading. The node measures the round-trip time for this message exchange and, acting on the assumption of a stable delay, defines the time half of the round-trip time after sending the synchronization request as the time when the server measured its clock reading. To increase the accuracy, this message exchange is done several times and the trial with the lowest delay is chosen. A drawback of this protocol is that several messages have to be exchanged between each node and the server, which will increase the network load and the energy consumption. Additionally, the medium access delay can vary extensively, hence, the stability assumption might not be accurate.

To exclude the uncertainty introduced by the medium access delay another approach is discussed by Elson et al. [81]. They present Reference-Broadcast Synchronization (RBS), which exploits the broadcast nature of the wireless communication channel. All nodes in the broadcast region are able to receive a wireless packet at almost the same time. Hence, they use the arrival time of a synchronization packet to measure their local clock readings. Then they exchange their measurements to calculate the relative clock offset between each pair of nodes. RBS not only excludes the uncertainty of medium access delay, which can get very high under high network load, but also enables post-facto synchronization: with RBS the clock offset between each pair of nodes is known, therefore, assuming short-term clock stability, the occurrence time of an event can be determined even if the nodes were not already synchronized; the synchronization can be done shortly after the event occurred.

An approach with less communication effort is presented by Mock et al. [80]. There, all the nodes can synchronize their clocks to the arrival time of periodically transmitted synchronization packets and do not have to exchange their clock readings between each other. Only one message per synchronization round is sufficient to keep the nodes synchronized. This is achieved by not adjusting the local clocks but by keeping a “virtual” clock, which is not a real clock but rather describes the transformation of the local time to the server time. A drawback of this protocol is the high computational effort to maintain the “virtual” clock. Whenever a time value has to be used, a new “virtual” clock value has to be computed.

3.3 Bounded Clock Offset Synchronization

To maintain a very low power consumption, clock synchronization has to be done as energy-efficiently as possible. In the following a simple bounded clock offset synchronization protocol is discussed and evaluated. The requirements for the protocol are to synchronize any number of nodes within one communication region—meaning that all nodes are direct neighbors of the synchronization device—and that the protocol is as power-efficient as possible.

Two factors determine the energy efficiency of the synchronization approach: the amount of computational effort and the amount of wireless communication. The three protocols presented in Section 3.2.4 rely on a different number m of messages for synchronization. In NTP the number of messages increases proportional ($\mathcal{O}(n)$) with an increasing number n of nodes. In RBS the number of messages increases with $\mathcal{O}(n^2)$, because each node has to communicate with each other instead of just with the server. In the concept by Mock et al. [80] only one message has to be broadcast to all nodes, no exchange between nodes is necessary. Hence, the number of messages also increases with an increasing number of nodes ($\mathcal{O}(n)$) but in

comparison with NTP only half of the messages are needed to synchronize all nodes.

Hence, a concept like Mock et al.'s [80] would consume less power than the other concepts due to less communication. But a drawback is that it has to be known when the server is transmitting its synchronization packets. Otherwise, the nodes would not be able to adapt their transceiver duty cycle to it and they would be forced to keep their transceivers active all the time to be able to receive the synchronization packets. This becomes especially a problem if a node is in an unsynchronized state, which can happen due to communication failures or when the node is powered up for the first time.

The novel idea is to use a synchronization protocol similar as the protocol described by Mock et al. but instead of using normal packets wake-up messages are directly used for synchronization. The advantage over the protocol described by Mock et al. is that the main transceiver can be turned off all the time. No accurate duty cycling timers have to be maintained and it becomes possible to send synchronization messages at any time. No predetermined time window has to be met to send synchronization messages, because the nodes are able to receive them at any time. Additionally, to further simplify the protocol no "virtual" clocks are maintained but a limit for the maximum allowed clock offset is introduced.

In the following the factors which influence the clock offset are investigated. Furthermore, a jitter analysis for a specific sensor node platform is performed to determine its synchronization accuracy.

3.3.1 Keeping the Clock Offset Bounded

Most applications do not need a perfectly synchronized local clock. It is sufficient to limit the maximum allowed clock offset to ΔT_{lim} . This limit is determined by two factors: the maximum synchronization error ϵ_{max} and the clock offset between two synchronization points with worst-case assumptions for the clock frequency. The maximum allowed clock offset can be described by: $\Delta T_{\text{lim}} = \epsilon_{\text{max}} + \max \Delta T$.

A typical 32 kHz clock has a frequency tolerance of about 20 ppm. This value must be increased by 15 ppm assuming a maximum temperature deviation of 20°C from its nominal temperature. Hence, the maximum deviation of f_0 will be around 35 ppm from which follows: $\max \Delta T = 35 \cdot 10^{-6} \delta_s$.

As outlined in Section 3.2.3 several delay sources influence the synchronization error ϵ . The influence of the medium access delay δ_{mac} can be excluded because the time of arrival of the wireless packet is used as the common time basis for synchronization instead of the point in time when the medium access process was started. As already stated above, the propagation delay can be excluded as well.

The jitter of the transmission delay will have the main impact on the synchronization error. The jitter is predominantly induced by the oscilla-

tor settling time of the transceiver and by the asynchronism between the sampling times in the receiver and the corresponding bit edges.

The operating delay δ_o strongly depends on how the protocol is implemented. For instance, if a real time operating system is used and the wireless packet is processed in a task then the delay is typically in the region of up to 1 ms, but not predictable. To mitigate this effect, the processing of the synchronization packet has to be decoupled from measuring its time of arrival. Then, regardless of the actual time when the processing is done, the time of arrival can be used for synchronization. To decouple the timing of these events a timestamp for the wireless packet has to be created as soon as possible. Depending on the system two solutions are possible.

First, the timestamp can be created within the interrupt service routine (ISR) of the microcontroller. In the ISR all interrupts are processed. This is done independently from the operating system task scheduler and therefore the maximum delay until the ISR becomes active is much shorter. Usually, the ISR will become active when the microcontroller is woken up or when an already running ISR event has finished. Typically, only very short functions are allowed in an ISR which assures short ISR processing times. Therefore, the time it takes a microcontroller to wake up has a greater impact than the ISR processing time. Hence, this wake-up time mostly determines the maximum delay for an ISR to become active.

Second, a hardware timer with capture capability can be used. Such a hardware timer can run independently from the core of the microcontroller and can create a timestamp instantaneously when an interrupt is triggered by the transceiver. Thus, to create a timestamp the microcontroller does not have to be woken up. The ISR has to become active only to associate the already created timestamp to the received packet. Using a hardware timer is the most accurate solution to create timestamps but relies on hardware functionality and is not platform independent.

3.3.2 Jitter Analysis

The Bounded Clock Offset Synchronization is evaluated for two cases: using a duty-cycled transceiver and using a wake-up receiver. For all measurements sensor nodes, as depicted in Figure 3.6, where used which are equipped with a EFM32LG332 [92] microcontroller, an AT86RF212 [14] transceiver, a wake-up receiver [39], and an OV-7604-C7 32 kHz clock [75].

As mentioned in Section 3.2.1, a clock in combination with a microcontroller is typically used to provide the local time for a wireless sensor node. The precision of the local time clearly depends on the precision of the clock, in this case the 32 kHz clock OV-7604-C7. Additionally, when obtaining a reference time wirelessly, several elements in this communication chain are influenced by this low-frequency clock as well. For instance, to save energy the peripherals of the microcontroller are driven by this low-frequency clock

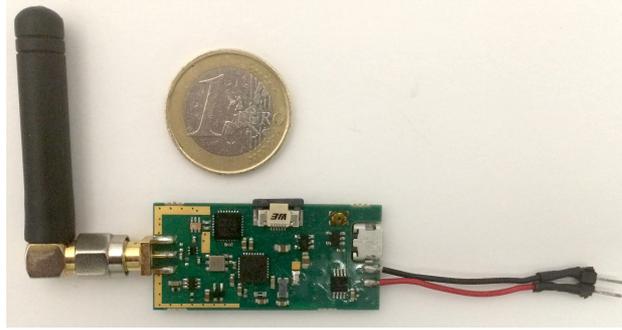


Figure 3.6: Airbus Group Innovations wireless sensor node with WUR.

instead of a high-frequency clock. One of these peripherals are low-frequency timers, which are used to capture the time when the transceiver, respectively WUR, received a synchronization packet. Another element influenced in the communication chain is the WUR itself: the examined wireless sensor node uses the same 32 kHz clock to drive the peripherals of the microcontroller as well as driving the WUR.

First the quality of the clock is analyzed by measuring the jitter of its pulses with an oscilloscope. The histogram of this jitter is depicted in Figure 3.7.

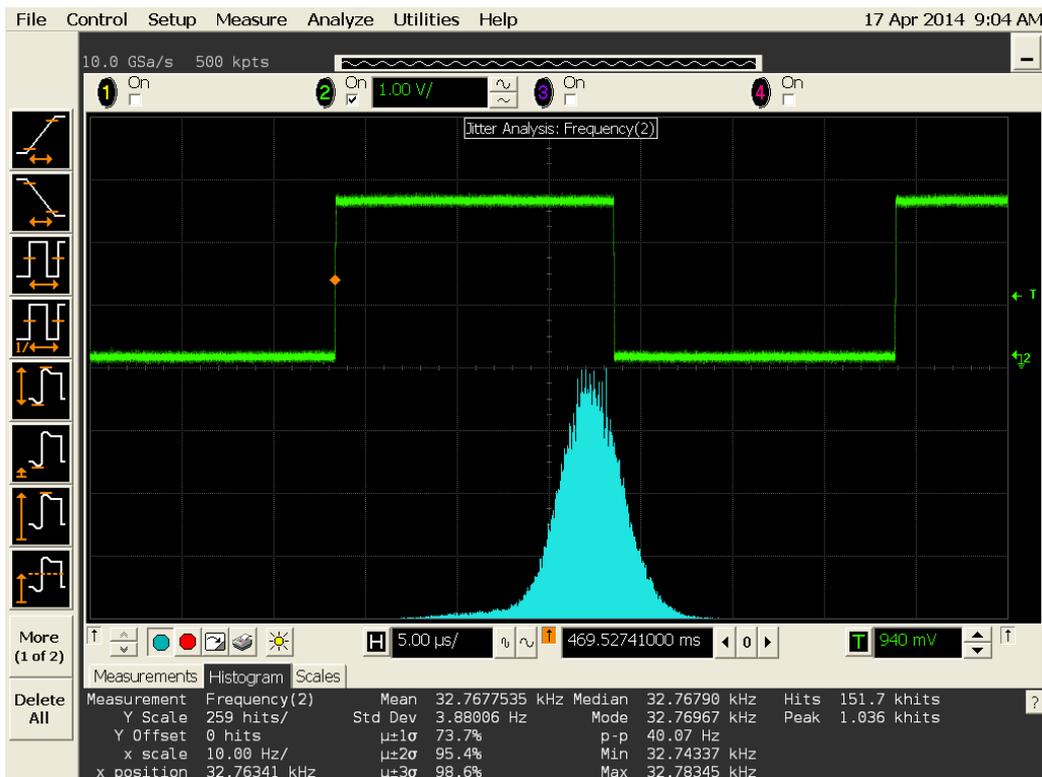


Figure 3.7: Jitter of an OV-7604-C7 32 kHz clock [75].

The jitter of the clock is measured more than 150,000 times to create its

histogram. As expected, it follows a Gaussian distribution with its mean at the nominal frequency of 32.768 kHz. Its standard deviation is measured as $\sigma = 3.88$ Hz, from which follows a deviation of around 3.7 ns for the duration of one clock pulse.

The 32 kHz clock is not the only source of imprecision when calculating the communication delay. As mentioned in Section 3.3.1, the relevant sources are the jitter of the transmission delay δ_t and the jitter of the operating delay δ_o . All other delay sources in the communication chain can be excluded or neglected.

To measure the magnitude of the jitter, three points in time can be distinguished: the time IRQ_{tx} when the transmitting node raises an interrupt to indicate to the transceiver to start the transmission of the synchronization packet, the time RF_{tx} when the PLL of the transceiver is settled and the actual transmission starts (respectively RF_{wur} when a wake-up message is transmitted instead of an IEEE802.15.4 data packet), and the time IRQ_{rx} when the receiving node recognizes the interrupt of an arriving synchronization packet within its ISR. The time IRQ_{tx} is measured by identifying the rising edges of the interrupt signals at the appropriate microcontroller PIN. The time $\text{RF}_{\text{tx}} / \text{RF}_{\text{wur}}$ is measured by identifying the time when the magnitude of the radio signal is the first time above a certain threshold (the antenna connectors of two nodes were connected to each other and to an oscilloscope by cables via a multiplexer). The time IRQ_{rx} is measured by toggling the status of a microcontroller PIN within the ISR and measuring its edge.

First, the jitter of the delay between indicating the start of the transmission and actually starting the transmission is evaluated. In Figure 3.8 the histogram of this jitter is depicted for (a) when transmitting an IEEE802.15.4 data packet and (b) when transmitting a wake-up message. Both cases have to be distinguished, because the transceiver uses a different transmission mode to transmit a wake-up message.

On average, it takes 50.28 μs to start the transmission of a IEEE802.15.4 data packet and 43.86 μs to start the transmission of an wake-up message. These different numbers have to be taken into account when calculating the total communication delay. The jitter of this delay is around ± 0.55 μs for the data packet and ± 0.63 μs for the wake-up message.

Second, the jitter of the total communication delay is evaluated: the delay between IRQ_{tx} and IRQ_{rx} . The jitter is mainly influenced by the sampling frequency of the transceiver. The radio wave arrives anywhere between two sampling points, hence, the delay varies at least with the sampling period.

Data rates of 20 kb/s, 40 kb/s, and 250 kb/s are used, which corresponds to sample rates of 300 kchips/s, 600 kchips/s, and 1000 kchips/s, respectively, as stated in the transceiver data sheet [14]. The histogram of the jitter is depicted in Figure 3.9 for a data rate of 20 kb/s and of 250 kb/s. As can be seen, the maximum deviation is slightly higher than two times the sample

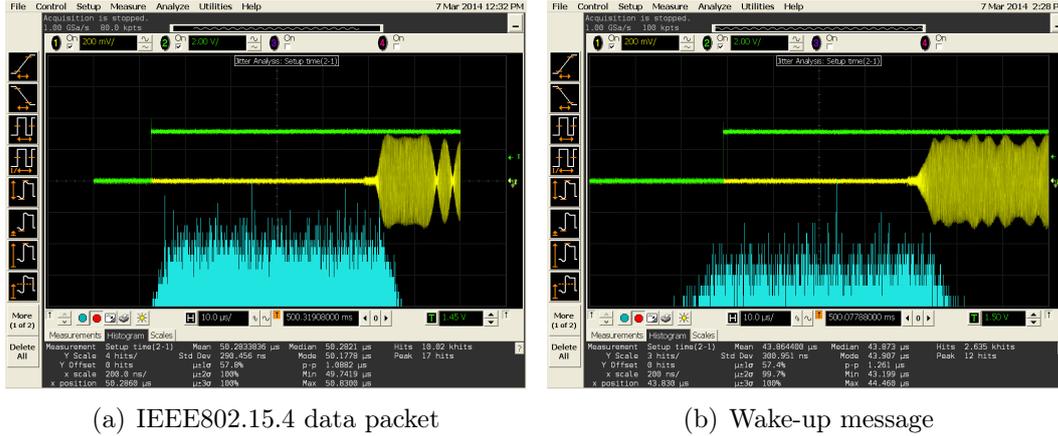


Figure 3.8: Jitter of the delay between IRQ_{tx} and (a) RF_{tx} or (b) RF_{wurr} .

period. A similar characteristic for the jitter can be seen for a data rate of 40 kb/s.

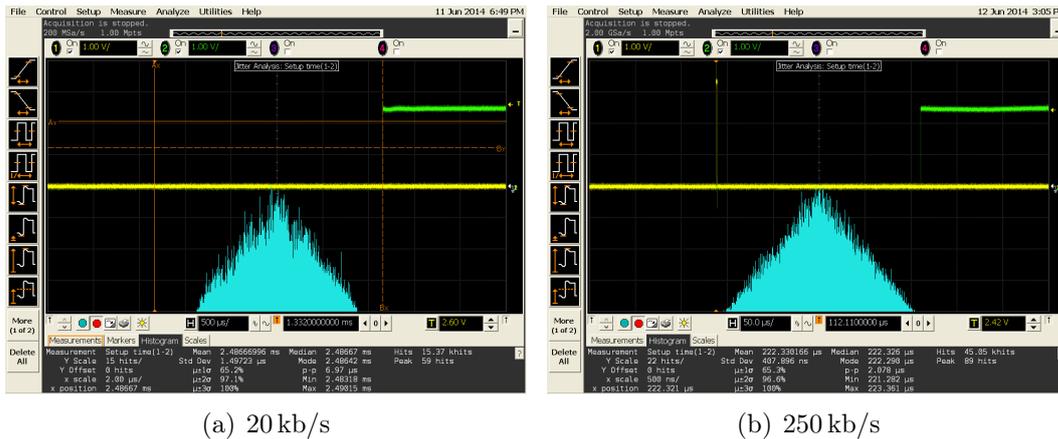


Figure 3.9: Jitter of the delay between IRQ_{tx} and IRQ_{rx} using an AT86RF212 transceiver [14].

The same analysis is done using a wake-up receiver. The greatest difference is that a wake-up receiver has a significantly lower data rate than a “normal” transceiver, which results in a higher synchronization error, as can be seen in Figure 3.10. The histogram follows a uniform distribution and has a range as broad as the sample period: around $244 \mu\text{s}$ at a sample rate of 4096 samples/s in Figure 3.10(a) and around $31 \mu\text{s}$ at a sample rate of 32768 samples/s in Figure 3.10(b).

Table 3.1 shows the characteristics of the jitter for different transceiver and wake-up receiver data rates. The maximum variation decreases proportional to the decrease of the sample period, hence, the synchronization error

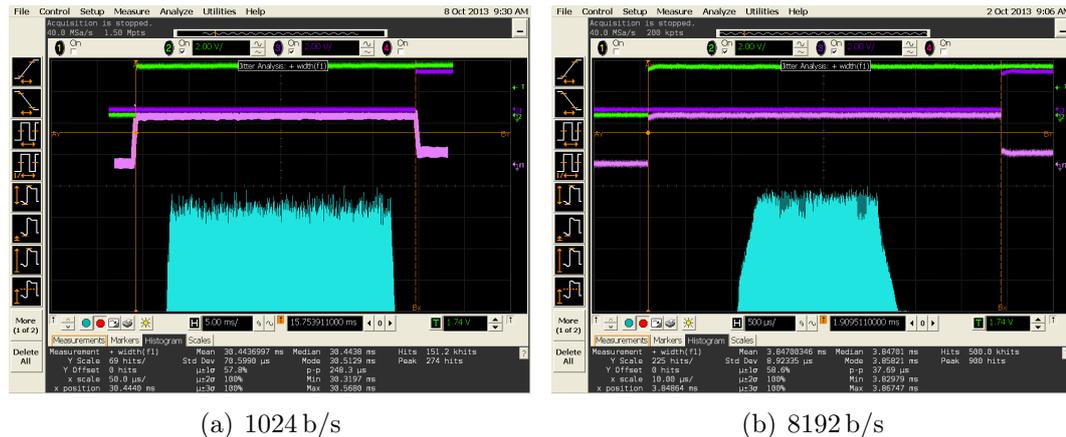


Figure 3.10: Jitter of the delay between IRQ_{tx} and IRQ_{rx} using the wake-up receiver design from Milosiu et al. [39].

decreases with an increasing data rate.

Table 3.1: Jitter characteristic for varying data rates

	data rate	sample period	σ	max – min
AT86RF212 [14]	20 kb/s	3.33 μ s	1.48 μ s	7.08 μ s
	40 kb/s	1.67 μ s	0.78 μ s	3.98 μ s
	250 kb/s	1.00 μ s	0.41 μ s	2.08 μ s
wake-up receiver [39]	1 kb/s	244.14 μ s	70.60 μ s	248.3 μ s
	8 kb/s	30.52 μ s	8.92 μ s	37.7 μ s

3.4 Evaluation

First, the power consumption for synchronization when using a duty-cycled transceiver is evaluated. This represents the baseline the wake-up receiver based synchronization protocol will be compared to and corresponds to the synchronization protocol of Mock et al. when nodes are already in a synchronized state. Second, the power consumption for synchronization when using wake-up receivers is analyzed and both approaches are compared to each other.

3.4.1 Using a duty-cycled Transceiver

The synchronization message when using a duty-cycled transceiver has to be sent at certain points in time. Only then will it be possible for the nodes to

match their transceiver ON times with the time when the synchronization message will be transmitted. Hence, it has to be guaranteed that it is possible to access the medium at these points in time.

The necessary ON time is determined by the wake-up delay, the medium access delay, the transmission duration, and the maximum clock drift. According to Equation (2.3), the mean power consumption of a state-of-the-art transceiver, for instance the IEEE802.15.4 capable AT86RF212 [14], is around $400 \mu\text{W}$ when using a synchronization period of 1 s at a data rate of 20 kb/s. Different data rates result in a different necessary ON time and therefore in a different mean power consumption. This behavior can be seen in Figure 3.11 where the mean power consumption for varying synchronization periods is shown. When using a synchronization period greater than 100 s all curves approach the same asymptote at around $2.8 \mu\text{W}$ because the difference in the varying ON time is significantly smaller than the much longer OFF time. Additionally, the behavior of a theoretical transceiver with similar power consumption values and settling times but with a much higher data rate of 100 Mb/s is depicted. This theoretical transceiver shows that there is a lower bound for the mean power consumption.

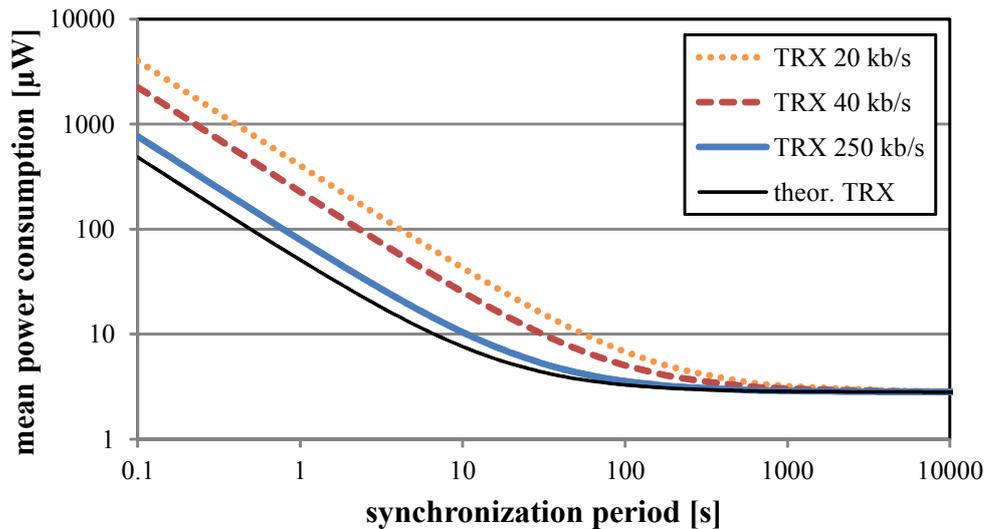


Figure 3.11: Mean power consumption of an AT86RF212 [14] transceiver for varying synchronization periods, together with a theoretical transceiver with very high data rate.

It is a great effort to implement a synchronized duty cycle for a transceiver: a node has to maintain an accurate sleep and wake-up timer for the transceiver, preferably implemented in hardware. Additionally, a recovery process has to be implemented for synchronizing for the first time or when synchronization was lost due to communication failures; a solution would be, for example,

to turn on the transceiver all the time until a first synchronization packet was received. When using a low duty cycle this time can be quite long which results in a high power consumption. Additionally, it is necessary to schedule synchronization messages in order to meet the listening times of duty-cycled transceivers. Hence, a more complex medium access protocol has to be used, a simple CSMA protocol does not support guaranteed sending times. To meet these challenges a wake-up receiver can be used, because then no specific ON time has to be met; a synchronization packet can be sent at any time.

3.4.2 Using a Wake-up Receiver

The power consumption of a WUR is very low, however, the sensitivity and data rate of a wake-up receiver is usually lower, as well. For instance, the wake-up receiver design by Milosiu et al. [39] has a sensitivity of -80 dBm, independent from the data rate. An AT86RF212 has a sensitivity of -110 dBm at 20 kb/s, -108 dBm at 40 kb/s, and -101 dBm at 250 kb/s. This increased sensitivity corresponds to a greater communication region, but as shown in [7] for aeronautical use-cases a sensitivity of -80 dBm is sufficient to support a communication region of about 10 m. Therefore, this lower sensitivity has no drawback for the aeronautical use case but might have one in other scenarios.

When using wake-up messages for synchronization, no accurate duty cycling timers have to be maintained and no power-hungry recovery process has to be implemented. Furthermore, it is possible to send synchronization messages at any time. Additionally, without the need to keep a low duty cycle it is feasible to send a synchronization packet once a second without increasing the power consumption significantly. This hardens the synchronization protocol against communication failures compared to using a transceiver with a low duty cycle, where for instance only every 100 seconds a synchronization packet is transmitted.

The drawback when using a wake-up receiver is its high synchronization inaccuracy due to its low data rate. Using a main transceiver results in a much higher synchronization accuracy. However, because of the long synchronization period when using a main transceiver the local clocks will drift further away and therefore the resulting clock offset will be in the same region as when using wake-up receivers.

The mean power consumption while keeping a maximum clock offset is investigated for both synchronization strategies: when using a duty-cycled transceiver and when using a wake-up receiver. The synchronization period is varied for the duty-cycled transceiver while a constant period of 1 s is kept for the wake-up receiver because its energy consumption is nearly independent from the synchronization period. As can be seen in Figure 3.12, at a data rate of 20 kb/s or 40 kb/s the mean power consumption of a duty-cycled transceiver is always higher than that of a wake-up receiver. When

using a data rate of 250 kb/s the mean power consumption of a duty-cycled transceiver is only smaller when limiting the clock offset to $< 240 \mu\text{s}$, which corresponds to a synchronization period of $< 7\text{s}$. Additionally, the behavior of a theoretical transceiver with a data rate of 100 Mb/s is calculated. Even if this theoretical transceiver with such a high data rate would exist, using a wake-up receiver still consumes less power when allowing a clock offset of 1 ms or greater.

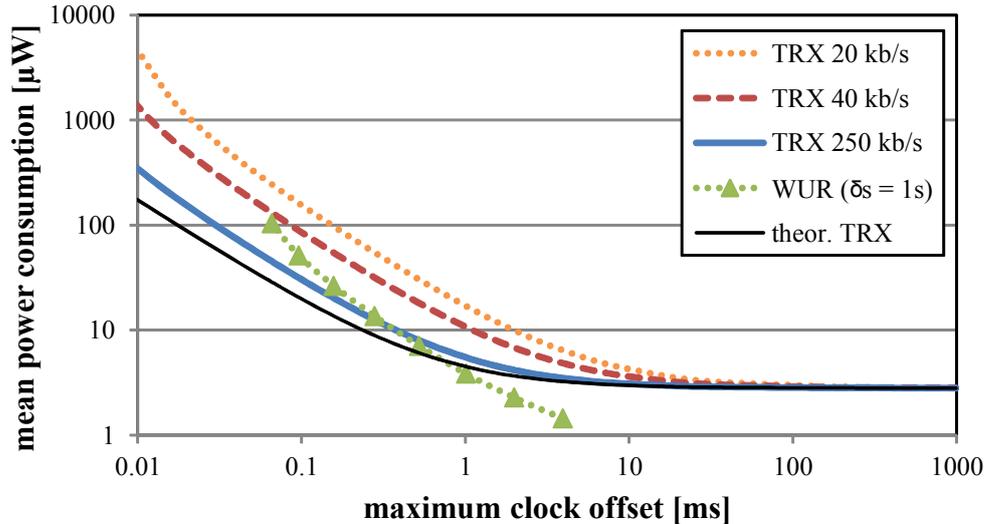


Figure 3.12: Comparison of the mean power consumption while keeping a maximum clock offset for a duty-cycled transceiver [14] and for a wake-up receiver [39].

Using a wake-up receiver means to integrate additional hardware which will increase the overall system costs. But using a wake-up receiver for synchronization is only one application scenario. Many other scenarios are possible which can decrease the overall system costs correspondingly. Comparing the additional costs with the costs for a GPS receiver—which can be used to get a (high precision) time reference—a wake-up receiver has still a lower power consumption. A state-of-the-art GPS receiver [76] consumes about 47 mW when searching for satellites and has a mean power consumption of about 8 mW when enough satellites are in range and once per second a new location / time update is done. These values are still magnitudes higher compared to the consumption of a wake-up receiver. Additionally, a GPS receiver can rarely be used indoor, whereas a wake-up receiver system is location independent.

3.5 Conclusion

One of the prerequisites to use WSNs for flight applications is the need to create a common time basis very energy-efficiently. It is feasible to use a wake-up receiver for creating this common time base by using wake-up messages to synchronize the local time of the sensor nodes. In doing so, the power consumption for synchronization is reduced while simultaneously increasing the robustness of the synchronization process against packet loss. Additionally, it is not necessary to have a complex medium access protocol; no extra scheduling effort has to be spent.

Only a duty-cycled transceiver will reach very high synchronization accuracy but at the cost of a high power consumption. If very low power consumption is required, wake-up receivers are the better choice. For instance, if allowing 1 ms maximum clock offset, using a wake-up receiver reduces the mean power consumption by 28 % compared to using a duty-cycled transceiver at the same maximum allowed clock offset. Additionally, the impact of a lost synchronization packet is much smaller due to the shorter synchronization period. A duty-cycled transceiver can only receive a packet each 30 s, while a wake-up receiver can receive synchronization packets any time, e.g. once a second. If a duty-cycled receiver with the same mean power consumption is used then only a maximum clock offset of 2.48 ms could be guaranteed and only each 70 s a synchronization packet would be transmitted.

4 Clustering using Wake-up Receivers

To avoid redundant transmissions nodes with similar sensor readings are clustered together and only their cluster leader will transmit its sensor readings as the cluster representative. Building the clusters is done in a very energy-efficient way by utilizing wake-up receivers.

4.1 Introduction

Many applications in the aeronautical domain require periodic transmissions of sensor readings. Especially in the SHM domain a lot of sensor readings are generated and the transmission of these will consume a lot of energy. Because of the inaccessibility of SHM wireless sensor nodes, energy-efficient operation is the key requirement for this application domain. One approach to reduce energy consumption is to group nodes together that are generating similar sensor readings. Then, only one sensor reading out of a cluster is sufficient to approximate the readings for the whole cluster; redundant transmissions can be avoided and, thus, the mean energy consumption is reduced.

Much research has been undertaken in the last years to decrease energy consumption by clustering [93–106]. For example, Tulone et al. [97] present the clustering protocol Probabilistic Adaptable Query system (PAQ). PAQ divides the nodes into clusters in which each node generates similar sensor readings. The cluster leader then uses a time series forecasting model and sends the model coefficients together with the cluster membership lists to the sink. Because of the data similarity within a cluster, the sink can predict the sensor readings for all cluster members. But PAQ has the drawback of a high energy consumption for the clustering process itself. If re-clustering is often necessary, more energy will be spent for the clustering than can be saved afterward.

The novel idea is to use a clustering protocol similar to PAQ but with a more energy-efficient clustering process; especially for a high number of nodes. This is accomplished by utilizing wake-up receivers during clustering: information about sensor readings is disseminated by wake-up receiver messages instead of real data packets, which makes it possible to turn off the main transceiver most of the time.

This WUR-based clustering protocol is evaluated on the basis of data from

a real-life measurement campaign as well as a generic data model. The performance of the protocol is evaluated in terms of energy consumption and data throughput by exhaustive parameter variation. Additionally, this protocol is implemented on wireless sensor nodes to verify the simulation results and to demonstrate the functionality of the protocol in real life.

4.2 Background

First, background information is given about sensor readings and how to map these readings to wake-up messages before the clustering protocol is presented in more detail.

4.2.1 Sensor Readings

To reduce redundant data transmissions, nodes which have similar sensor readings will be clustered together. Therefore, the co-domain of all sensor readings is divided into intervals of width δ (for more detail, see Section 4.2.2). All sensor readings within the same interval are considered similar. Hence, for a group V of sensor readings follows:

$$\forall v_1, v_2 \in V : |v_1 - v_2| < \delta$$

The generation and stability of clusters strongly depends on the similarity of sensor readings and their development. To inspect this behavior two sets of sensor readings are considered in the simulations; a set of real sensor measurements and an artificial data model for flexibility, which aims for modeling the macroscopical behavior of the real sensor measurements.

Real Measurements

A set of real measurements is obtained from the EU project “Ideal Cabin Environment (ICE)” [107]. In this project the Fraunhofer Institute for Building Physics has constructed a flight laboratory which consists of a front aircraft segment (A310-200) installed into a low-pressure climate chamber. This climate chamber can achieve realistic temperature profiles inside the cabin by cooling or heating the complete outer fuselage of the aircraft (see [107]). In the cabin of this aircraft segment several temperature sensors are installed. The obtained measurement samples are one-minute mean values. These samples are linearly interpolated to extend the measurement set to one-second samples. These sensor readings are shown in Figure 4.1.

As can be seen, readings of groups of sensor nodes are showing a common behavior. They follow the same longterm drift and additionally fluctuate individually in a lower scale. In order to avoid optimizing the protocol for one specific set of sensor readings an artificial data model is established which

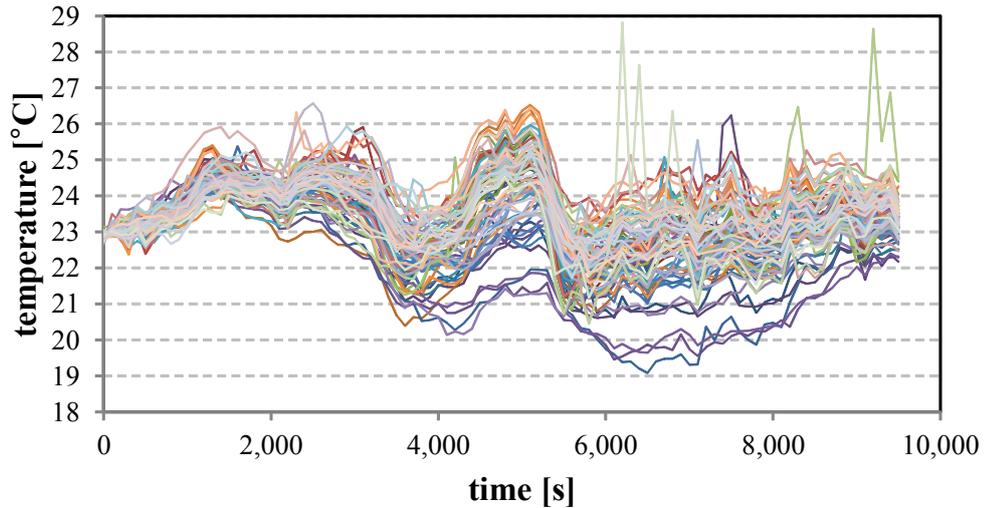


Figure 4.1: Sensor readings from the EU project “ICE: Ideal Cabin Environment Project” [107].

aims for modeling this behavior. With an artificial data model, it becomes possible to vary sensor readings and, thus, to increase the significance of the results.

Data Model

Two main behaviors are abstracted from the set of real sensor readings; the individual drift of readings of sensor nodes and the common drift of readings of groups of sensor nodes. To model this behavior, a sensor reading is composed of two mean values generated by Markov chains and an additional random Gaussian white noise with zero mean and $\sigma^2 = 0.01$.

The first Markov chain shown in Figure 4.2 represents the individual drift of a node. It consists of five linearly connected states where each state represents a fixed mean value. The middle state represents a mean value of 0 and the consecutive states each represent additionally $\pm 0.1^\circ\text{C}$. Hence, the maximum additional mean equals $\pm 0.2^\circ\text{C}$, which represents approximately the individual drift observed in the set of real measurements. Each state has a probability of 0.9 to stay in the state. The probability for transitioning from the middle state to one of the adjacent states is 0.05 each. This transition probability decreases linearly to the outer states until it reaches zero.

The same linear transition behavior is valid for the second Markov chain which models the common group behavior. This Markov chain is depicted in Figure 4.3 and consist of seven states which represent mean values separated by 0.5°C . These mean values vary from 20°C to 23°C . This range approximately represents the mean range the sensor readings of the real mea-

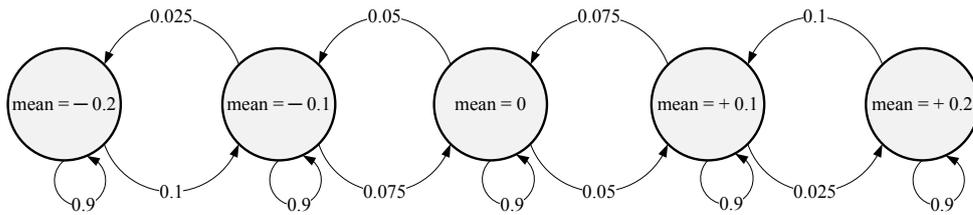


Figure 4.2: First Markov chain which models the individual drift for each node.

measurements are in. Nodes which are currently following the same longterm drift are the nodes which are in the same state of the second Markov chain. Hence, nodes will change from one group to another if the state of their second Markov chain changes.

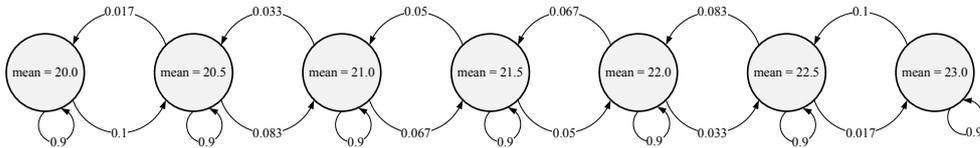


Figure 4.3: Second Markov chain which models the common group behavior for each node.

In Figure 4.4 sensor readings of four nodes are shown that were generated using the Markov chains. As can be seen, in this example, node 1 and 2 respectively node 3 and 4 show a common group behavior.

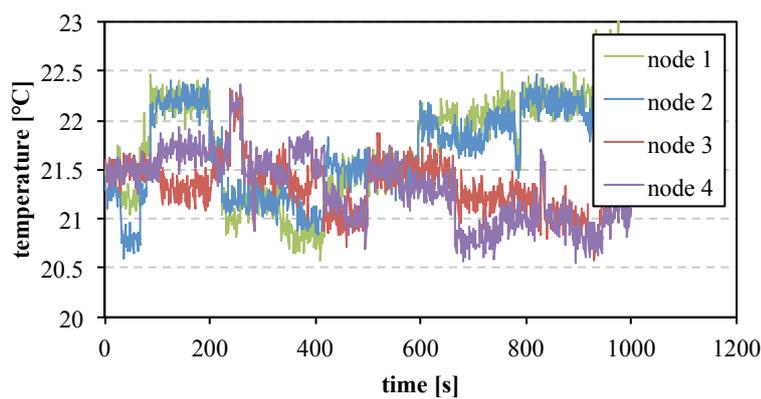


Figure 4.4: Sensor readings of four nodes generated by the data model.

4.2.2 Mapping Measurements to Wake-up Messages

The idea of the new clustering protocol is to disseminate information about sensor readings by using wake-up messages instead of normal data packets. In order to do this, the large set of sensor readings has to be mapped to a small set of available wake-up messages. As a first solution, this is done by dividing the co-domain of the sensor readings in fixed, pre-defined intervals and each interval will be associated with one specific wake-up message.

The cardinality n of the set of wake-up messages thereby defines the possible mapping co-domain or the size of the intervals, respectively. The interval width is δ . All measured values have to be mapped to such an interval. Thus, the mapping corresponds to a quantizing problem and the width of one bin corresponds to δ . All measurements within one bin are similar to each other or in other words the difference of each other is smaller than δ .

With the relatively small cardinality of the wake-up message set, there is a trade-off between possible values of δ and the mapping co-domain Δ , which can be supported without clipping. For example, if there are 100 messages available then 100 bins are possible. If $\delta = 0.5$ is desired, the maximum mapping co-domain is $\Delta = 100\delta = 50$. Meaning, for example when temperature is measured, values between -10°C and $+40^\circ\text{C}$ can be correctly mapped to a pattern. Hence, if the measuring co-domain of the measurements differ from the mapping co-domain, the mapping may not fit. To support all measurements, even if they are outside the mapping co-domain, all measurements outside it have to be mapped to the nearest bin, despite the fact that they are not similar to the other measurements within this bin. Alternatively, if a bigger value for δ is chosen, the mapping co-domain will increase and more values can be mapped without clipping. The calculation of the co-domain is done by $\Delta = n\delta$.

4.3 Protocol Design

To reduce data transmissions nodes will be clustered into groups with similar sensor readings. The clusters are independent of the topology of the network with the constraint that the sink and each node is in communication range of all others. The only parameter for clustering nodes together is the similarity of their sensor readings. Hence, it is possible that nodes which are installed at different places are nevertheless clustered together because of their similar sensor readings.

After the clusters are established, a representative of each group k will transmit its readings as the “cluster reading” c_k . All other cluster members will not transmit their sensor readings. The sink will approximate all sensor readings of a cluster by its corresponding “cluster reading”. Hence, redundant data transmissions can be reduced and energy can be saved.

The clustering process has to be done energy-efficiently because otherwise, if the clusters are unstable, more energy might be spent for the clustering process itself than can be saved afterward. For the medium access a fixed, known TDMA schedule is used. Therefore, each node knows implicitly who is sending at any given time.

In the following section, first, the WUR-based clustering protocol is described in detail together with a conventional method. The conventional method is similar to PAQ, the clustering protocol by Tulone et al. [97], and acts as a baseline for comparison with the wake-up receiver approach. Second, a monitoring phase is introduced in which the validity of the clusters is monitored while energy is saved by reducing redundant data transmissions. Third, an optional protocol extension is presented which supports the transmission of measurements generated during the clustering process, which otherwise would not have been transmitted. Without this extension the sink will have no knowledge about these measurement values.

4.3.1 Clustering

The novel concept of the WUR-based clustering protocol is to use wake-up messages for disseminating information about sensor readings in a very energy-efficient way.

To cluster nodes which have similar sensor readings, first, each node has to have information about the sensor readings of all other nodes. Second, each node calculates the similarity of its readings to the readings of all other nodes. All nodes with similar readings will be clustered together and one node of each cluster will be chosen as its leader which will then transmit a *leader announcement* to inform the sink about the cluster and its members. This leader can only be chosen correctly if all cluster members are aware of each other.

This clustering process can be divided into two phases. Phase I is the information gathering mode where the information about sensor readings is broadcast. Therefore, the sensor readings are mapped to fixed, pre-defined intervals of width δ (see Section 4.2.2) and all readings within an interval are regarded as similar. Each node analyzes the similarity of its actual sensor reading to the broadcast ones. This is done m times to harden the protocol against outliers of the sensor readings or against communication failures. With this similarity information the clusters are established. Each node knows which other node has similar sensor readings and hence, knows its cluster members.

In Phase II, for each cluster the node with the smallest ID becomes the cluster leader and transmits a *leader announcement* message which includes the cluster membership list. After the clustering process is finished, each cluster leader will act as the cluster representative and will transmit its sensor readings to the sink.

Conventional Approach

A conventional approach without using a wake-up receiver can be seen in Figure 4.5. Phase I consists of m frames which are used to disseminate the sensor readings. For each frame each node transmits its actual reading at its own slot. In the other slots it listens for the readings of the other nodes. Therefore, the transceivers of all nodes have to be active for the entire Phase I. If more than *thold* similar readings from another node are received, this node is associated to the same cluster.

Phase II consists of only one frame in which the *leader announcements* are transmitted. The *leader announcement* message contains the leader ID and the cluster membership list. Before a leader transmits its announcement it listens to all prior slots to verify that no other node has chosen itself as the leader for this cluster and transmitted an announcement. If this was the case then the earlier announcing node becomes the actual leader and the other leader becomes a cluster member and does not transmit a *leader announcement*. All cluster members listen all the time to receive the announcements. They will check if they are correctly listed in the corresponding cluster membership list. If they are not correctly listed they will become cluster leaders themselves. In Phase II the transceivers of cluster members have to be active all the time; the transceivers of cluster leaders have to be active until the end of their own slots. The high energy consumption of the conventional clustering protocol results from the nearly always active transceiver. To lower the energy consumption of the clustering protocol the activity of the transceiver is the most promising leverage.

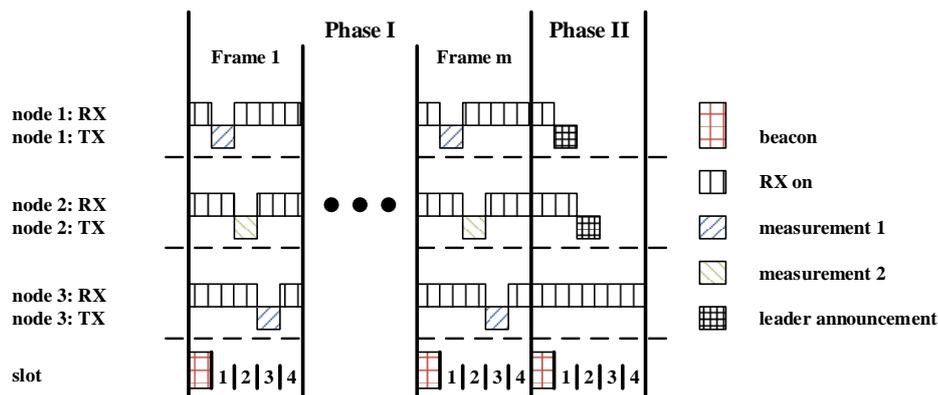


Figure 4.5: Example for the clustering protocol for three nodes **without** using a wake-up receiver.

Approach using Wake-Up Receiver

The novel idea to lower the energy consumption is to “mis-use” wake-up messages to disseminate information about measurements. The transceivers can be turned off most of the time which will lead to significantly decreased power consumption. The ID of a wake-up message, which is normally used to wake-up specific nodes, is now used to disseminate the information about a node’s sensor readings. Figure 4.6 shows an example of the clustering protocol for three nodes when using wake-up receivers. For each frame in Phase I each node transmits a wake-up message which corresponds to its own measurement. It is chosen by mapping the sensor reading to tabs (see Section 4.2.2).

A tab represents a fixed predefined data co-domain range of width δ . Each tab is represented by a specific wake-up message. Unlike the conventional approach, the transceiver is powered down the rest of the time. The wake-up receiver is tuned to wake up the microcontroller when a message is received which corresponds to the node’s own transmitted wake-up message. This means the wake-up receiver will only wake up the node when a “similar reading” was received. Because of the fixed TDMA schedule, a node knows implicitly which node transmitted the wake-up message. After Phase I each node knows which other nodes measured similar sensor readings but it had to spent much less energy to gather this information.

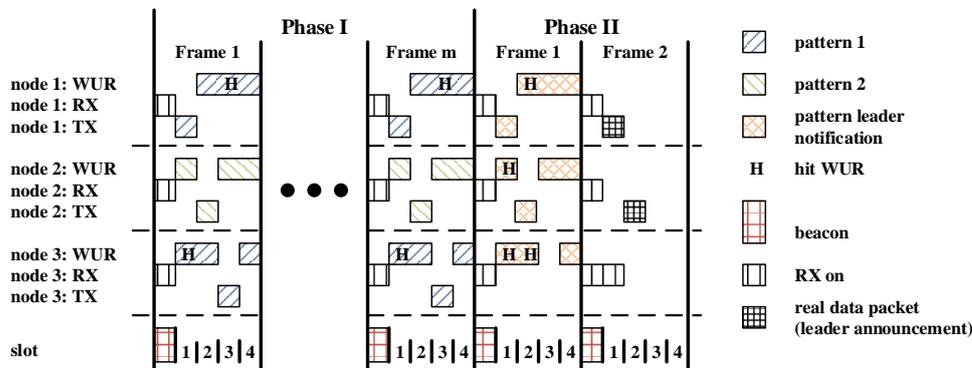


Figure 4.6: Example for the clustering protocol for three nodes when using wake-up receiver.

Phase II of the clustering protocol consists of two frames instead of one as in the conventional approach. The first frame is used to pre-announce which node will become a cluster leader. This pre-announcement is done per wake-up messages and, therefore, the transceiver can be powered down most of the time. Because of these pre-announcements the cluster members will know at which slot in the second frame they have to activate their transceivers to receive their *leader announcement*. Hence, the transceivers can be powered

down most of the time. Unlike the clustering protocol described in previous work [5] there is no third frame in Phase II because of a different approach for the monitoring phase.

4.3.2 Monitoring

The monitoring phase starts after the clustering process is finished. In the monitoring phase it is examined if the assumptions which led to the clusters are still valid.

Therefore, all cluster leaders will broadcast their sensor readings as “cluster readings”. If a “cluster reading” is dissimilar to the measurement of a cluster member, then, the cluster member will transmit an *outlier packet* to the sink which contains its measurement. This will ensure that the maximum approximation error at the sink is lower than δ . All monitoring methods have in common that if a node has to send more than i *outlier packets* within a time window t it will send a *reclustering request* to announce to the sink that its assumptions for the cluster are not valid any more. If the sink receives more than j *reclustering requests* it will initiate a reclustering. The minimum time before a reclustering can occur after the monitoring phase started equals i frames, because only after i frames the first *reclustering request* can be sent.

Regarding the volatility and trend components of the measured data, different methods are proposed to monitor if all cluster members still generate similar sensor readings. The first monitoring method performs well for moderately volatile data. Even if the data has trend components this monitoring method can perform well. The second monitoring method is well suited for the same data characteristics as the first method. Additionally, because of exponentially smoothing the data even more volatile data will give equivalent performance results. The third monitoring method is only very well suited for static data, but when well suited it will outperform the other methods regarding energy consumption.

First Monitoring Method

In the first monitoring method the cluster leaders will broadcast all their sensor readings—the “cluster readings”—and the cluster members will compare their sensor readings with them. If the own measurement is regarded as similar to the “cluster reading”, the cluster member will transmit nothing. At the end of each frame the sink received measurements of the cluster leader and possible *outlier packets* of their members. For these nodes the sink has exact values of their sensor readings. For all other nodes the sink knows implicitly that the difference of their sensor readings to the “cluster reading” is smaller than δ . Hence, the sink will approximate their measurements with the value of the corresponding “cluster reading”. The transceiver of cluster leaders only has to be active to transmit “cluster readings”. The transceiver

of cluster members has to be active to receive the measurement of their cluster leader. But because of the fixed TDMA schedule a cluster member knows when the leader will transmit its packet and, therefore, its transceiver has to be active for only one slot. Additionally, if an *outlier packet* has to be transmitted a member will be active for one more slot. If the member slot is scheduled before the leader slot the *outlier packet* will be transmitted in the next frame.

Second Monitoring Method

If the measured data is very volatile, the first monitoring method could lead to a lot of *outlier packets*. Therefore, in the second method the “cluster reading” is not equivalent to the real sensor reading of the cluster leader but is equivalent to an exponentially smoothed version of it. Exponentially smoothing a time series follows: $\hat{v}_n = \alpha\hat{v}_{n-1} + (1 - \alpha)v_n$, where $0 < \alpha < 1$.

The bigger α is, the smoother is the time series but the bigger is the offset between a real value and its smoothed version. Smoothing the measured data may lead to fewer *outlier packets* and in the long term to less reclustering but its effectiveness strongly depends on the measured data. To be able to exponentially smooth the real measurements of the cluster leader the cluster members and the sink have to remember the last smoothed value. Additionally, the cluster leader has to transmit a start value for the smoothing together with the *leader announcement*. The sink has to calculate the exponentially smoothed value as well, because this value becomes the new “cluster reading” and is used as the approximation for the values of the cluster members. Cluster leaders don’t transmit the exponentially smoothed versions of their values because then the sink knows exactly the values of cluster leaders instead of only their smoothed versions.

Third Monitoring Method

The last monitoring method is optimized for nearly constant sensor readings. In this method the “cluster reading” is fixed and does not vary over time and is determined once directly after the clustering process. Therefore, no node, neither the cluster members nor the cluster leader, has to transmit any measurements apart from *outlier packets*; this leads to an even smaller energy consumption. The sink will approximate all values with the “cluster reading” regardless of the role of a node. The “cluster reading” is defined by the middle of the tab a cluster leader chose during the clustering process. The cluster leader will choose that tab most of its sensor readings lay in during the clustering process. As for the exponentially smoothed method, the value of the middle of the tab has to be transmitted together with the *leader announcement*. This monitoring method is very energy-efficient for stable data, but for highly volatile data or data with strong trend components this

method will lead to a lot of *outlier packets* and in the long term to a lot of reclustering.

4.3.3 Protocol Extension

It is not possible to transmit all sensor readings generated during the clustering phase. In order to still have information about these values the protocol can be extended to transmit them later on.

During the monitoring phase the sink has knowledge about every sensor reading for all nodes, but during the clustering phase this is different. In the conventional approach in Phase I of the clustering real measurements are broadcast which means that the sink will have knowledge about these values. In Phase II only *leader announcements* are transmitted, which means the sink will have no information about measurements generated during this phase. Hence, each time a conventional reclustering is done the sink will have **one** measurement per node for which it has no knowledge.

In the novel approach in Phase I only wake-up messages are broadcast. Even if the sink is equipped with a wake-up receiver, it will not be able to receive all wake-up messages because a wake-up receiver, such as the one presented by Hambeck et al. [31], can only be tuned to one specific wake-up pattern at a time; the sink will have no knowledge about these measurements. During Phase II the sink will have no knowledge about measurements, as well. Hence, each time a WUR-based reclustering is done, the sink will have **five** measurements per node for which it has no knowledge (m measurements in Phase I and two measurements in Phase II).

The term goodput is used as the rate of measurements the sink has knowledge about. This includes all received measurements, received *outlier packets*, and approximated measurements.

$$\text{goodput} = \frac{\text{rcvd meas} + \text{outlier packets} + \text{approx}}{\text{total number of meas}} \quad (4.1)$$

The data transmission scheme can be extended to improve the possible goodput. With the extension nodes will transmit their measurements—which were not transmitted during the clustering—later on. This is possible because, most of the time, cluster members do not use their own slot during the monitoring phase. In doing so, the loss of knowledge can be prevented or at least alleviated. It is important not to delay the transmission of any other packet to maintain the protocol functionality and to guarantee data similarity. Therefore, a node will only transmit a postponed measurement if no *outlier packet* is present.

4.4 Evaluation

To simulate the whole chain of monitoring and reclustering, the protocol is implemented in OPNET network simulator. In the simulation scenario 77 nodes are equally distributed in a 5x5 m square and the sink is centrally positioned. This setup corresponds approximately to the setup in an aircraft and to the placement of the temperature sensors in the ICE project. The modeled wireless sensor node uses state-of-the-art components. The modeled components on this node are the microcontroller, the main transceiver and the wake-up receiver. The following devices are assumed: for the microcontroller the EFM32G210 [108] from Silicon Labs, for the main transceiver the AT86RF212 [14] from Atmel and the Wake-up Receiver designed by Hambeck et al. [31] within the EU project CHOSeN [69]. Table 4.1 shows the settings for the simulation.

Table 4.1: Simulation Settings

data rate	40 kbps	P_{tx}	39 mW
slot length	12.8 ms	P_{rx}	27.6 mW
frame length T	1 s	P_{wur}	0.0024 mW
slots per frame	77	P_{micro}	17.28 mW
beacon size	96 bit	m (frames in Phase I)	3
wake-up pattern size	64 bit	j (<i>reclustering requests</i>)	8
packet size	512 bit	t (time window size)	5.0 s
carrier frequency	868 MHz	α	0.9
transmit power	0 dBm	δ	0.5

In the simulated scenario each node is within communication range of each other and has knowledge about the fixed TDMA schedule. It knows implicitly who is sending at any given time. In the simulation each node generates a new sensor reading at the beginning of each frame. At first a perfect system without communication failures is assumed.

4.4.1 Clustering Phase

In this subsection only the clustering phase without the following monitoring phase is analyzed. Therefore, several simulation runs are performed using the artificial data model to test the clustering phase of the protocol under different situations. First, the total consumed energy per node during the clustering phase is analyzed, both using the WUR-based and the conventional clustering approach. Second, the scalability of both clustering protocols is compared. And last, the impact of incorrect wake-up message receptions is evaluated.

In Figure 4.7 the total consumed energy for Node 1 and Node 31 is shown when using the WUR-based approach. Both nodes are selected because the main power consumption of these nodes is separated by a distinct time period and, therefore, it is easier to distinguish between them. In this case, Node 1 acts as a cluster leader and Node 31 as its associated cluster member. The energy consumed by the wake-up receiver and the microcontroller is modeled and included into the total energy consumption but will not be visualized separately because their influence is marginal compared to the energy consumed by the main transceiver.

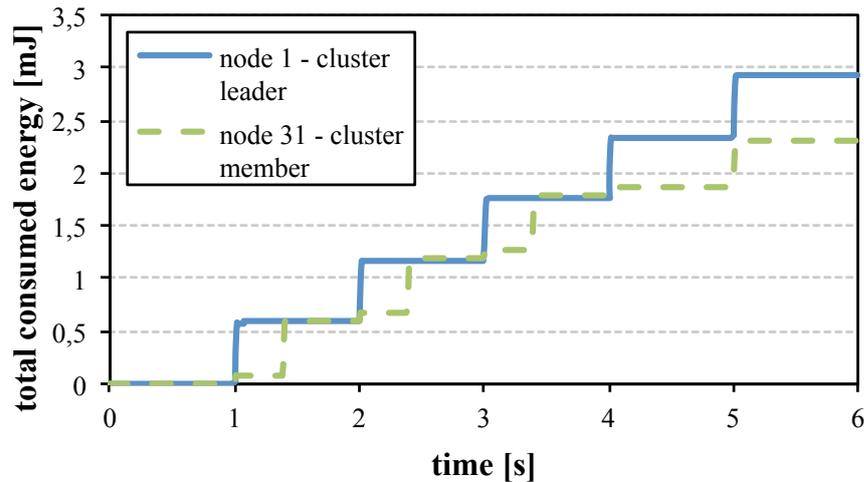


Figure 4.7: Total energy consumption of Node 1 and 31 while establishing the clusters for all 77 nodes with $m = 3$ and $\text{thold} = 2$ when using wake-up receiver.

The first frame starts with the arrival of the first beacon at simulation time $t = 1$ s. At that point in time the first peak of consumed energy is recognizable for Node 31. This peak occurs because the main transceiver is switched on for the reception of the beacon. Such a behavior can also be seen for each further beacon reception at $t = 2$ s, $t = 3$ s, $t = 4$ s, and $t = 6$ s. At time $t = 5$ s this small peak is not recognizable because the energy for listening at the directly following slot adds up to a big peak. Such a big peak can also be seen for each time the transceiver is active in Slot 31 where Node 31 is transmitting a wake-up message. This happens for example in the first frame at $t \approx 1.39$ s. These peaks are higher than the peaks receiving a beacon because the time to send a wake-up message is longer than that for receiving a beacon. A similar behavior can be seen for Node 1—the cluster leader. In the first three frames Node 1 consumes a similar amount of energy as Node 31 but at different slots. This behavior changes for the following frames: In Frame 4, Node 1 transmits a wake-up message to pre-announce its

status as a cluster leader which will be followed by a real packet—the *leader announcement*—in Frame 5. In contrast, Node 31 will keep its transceiver turned off during Frame 4 and will only enable it once in Frame 5 to receive the *leader announcement*.

For comparison, the total energy consumption for the conventional clustering approach is shown in Figure 4.8. As can be seen, the energy consumption of Node 1 and 31 is nearly equal. But now distinctive peaks are recognizable because the main transceiver is switched on most of the time. The clustering process is done within four frames instead of five but consumes approximately 110 mJ; significantly more than the corresponding 2.4 mJ for a cluster member in the WUR-based clustering approach, as shown in Figure 4.7.

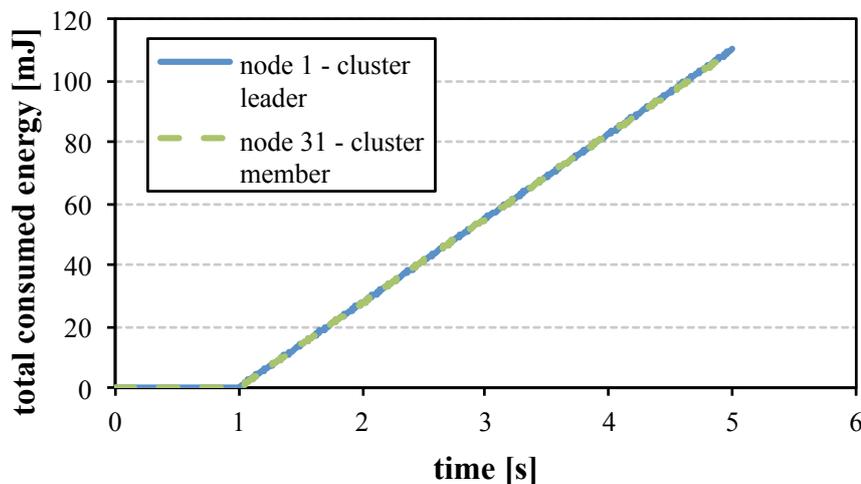


Figure 4.8: Total energy consumption of Node 1 and 31 while establishing the clusters for all 77 nodes with $m = 3$ and thold = 2 without using wake-up receiver.

Scalability

A major difference between both clustering approaches is that the energy consumption for the wake-up receiver based approach is almost independent of the number of nodes present. This situation changes for the conventional approach, which can be seen in Figure 4.9. The average energy consumption for the standard protocol increases linearly with an increasing number of nodes, while the graph for the WUR-based approach is independent from it and stays consistently around 2.4 mJ. This is the case because the main transceiver will not be switched on for more than one slot in any frame; either a wake-up message/packet will be transmitted or the transceiver will be active for one specific slot to receive a *leader announcement*. In the conventional

approach each node has to listen to all slots with their main transceiver which limits scalability. Conversely, the WUR-based clustering approach is highly scalable because an additional node does not increase the overall power need.

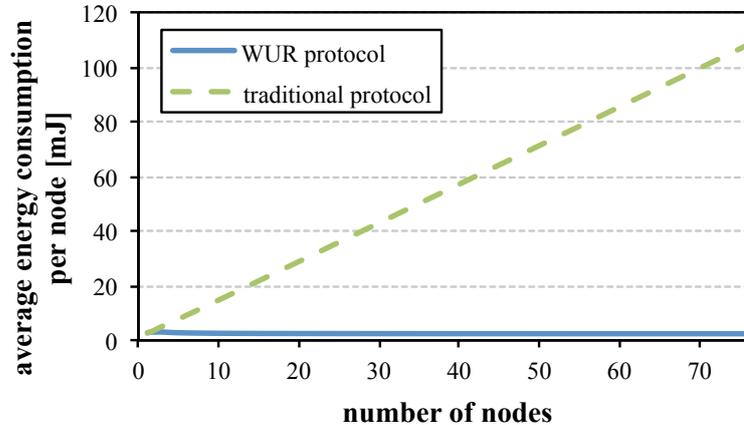


Figure 4.9: Systemwide average of energy consumption per node for the clustering phase over varying numbers of nodes.

Impact of wake-up receiver errors

In order to model the detection behavior of the Wake-up Receiver two error rates—the WUR errors—are introduced. The false wake-up rate P_{fwr} and the missed detection rate $\overline{P_{\text{det}}}$. For the missed detection rate holds $\overline{P_{\text{det}}} = 1 - P_{\text{det}}$, where P_{det} is the probability for a successful detection. The missed detection rate covers the situation when the wake-up receiver missed a wake-up message for which it was listening. On the other hand, the false wake-up rate covers the situation when the wake-up receiver announces a “correct” wake-up message although it received a “wrong” wake-up message it was not listening for.

To examine the impact of these WUR errors on the clustering phase, the false wake-up rate and the missed detection rate is alternately set to zero. Two different simulation sets are performed. First, the impact of a missed detection will be examined. Two kinds of wake-up messages can be missed, the wake-up message which disseminates the information about the sensor readings or the wake-up message which is used to pre-announce a *leader announcement*.

If a wake-up message in Phase I of the clustering is missed, a similar sensor reading will be treated as dissimilar. If more than $thold/2$ similar sensor readings of a node are falsely treated as dissimilar this node will not be associated to the same cluster as the receiving node. The consequence of such a situation greatly depends on the role of both nodes. If both nodes were

cluster members, there will be no impact at all. If the receiving node was the cluster leader, the transmitting node will not be included in its cluster member list and, hence, the transmitting node will become unnecessarily a cluster leader. Also, if the transmitting node was the cluster leader, the receiving node will either regard another node as the cluster leader or even regard itself as the new cluster leader. A missed *leader announcement* will force the node that missed it to become a cluster leader itself. Ultimately, both cases lead to more clusters but, nevertheless, the nodes within a cluster still have similar sensor readings. Hence, the number of members per cluster will decrease with an increasing number of leader and the overall power consumption increases.

On the other hand, if a false wake-up error occurs, a dissimilar sensor reading will be treated as similar. The impact of such an error is greater because it could result in a cluster leader treating a node with dissimilar sensor readings as a valid cluster member. Therefore, the number of members per cluster will increase with an increasing number of leaders. It can happen that a node will become a member in several clusters due to misinterpreted sensor readings. This means that the data sink may receive a cluster list in which members with dissimilar sensor readings are listed. Hence, a data prediction model will produce incorrect values for these nodes. This effect can be mitigated when data post-processing is done in the sink because the sink can recognize that a node is a member of several clusters and can mark its data as corrupt. In Phase II of the clustering protocol only one type of wake-up message will be sent, therefore, a false wake-up error cannot occur in Phase II.

As can be seen in Figure 4.10, three different curves are plotted; the number of leader before *announcement messages* were sent, the additional number of leader afterward and the total number of leader. For all curves confidence intervals with a confidence level of 95% are shown. The increased number of cluster leaders after the *announcement messages* is explained by nodes which did not find their ID on the member list and consequently became leader themselves.

What can be seen in Figure 4.10(a) is that—using the artificial data model—the number of leaders starts to change only around $\overline{P_{\text{det}}} \approx 0.01$. A lower missed detection rate did not have any influence on the clustering protocol. With an increasing rate of missed detections the number of leader increases as well until all nodes become a cluster leader. As can be seen in Figure 4.10(b), the false wake-up rate did not have any influence on the clustering protocol until $P_{\text{fwr}} \approx 0.05$. Only a very high error rate will have an impact on the functionality of the system which, under normal circumstances, will not occur.

The number m of sensor readings can be varied to adapt the protocol to the requirements of the specific applications. To reduce the impact of WUR errors, m can be increased, which consequently also increases the duration

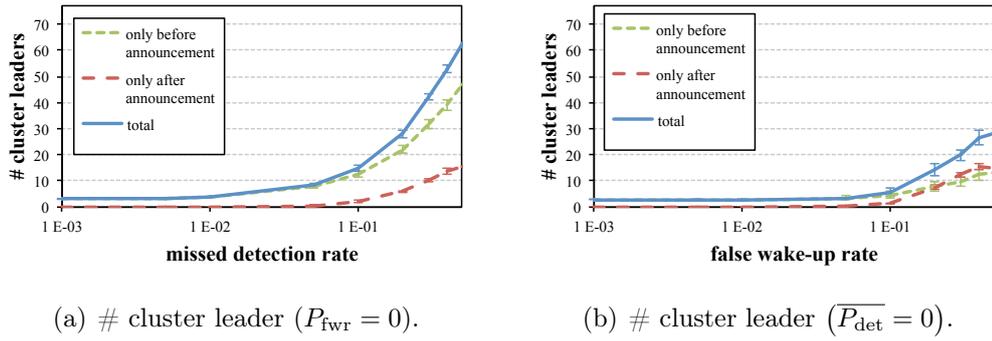


Figure 4.10: Number of cluster leader before an *leader announcement* was send, the additional leader afterward and the total number of cluster leader; Each with 95% confidence interval for $m = 3$, $thold = 2$ and $P_{\text{fwr}} = 0$.

of Phase I of the clustering process. E.g., with $m = 1$, the duration of Phase I is the smallest possible but WUR errors will instantaneously impact the performance of the clustering protocol, as can be seen in Figure 4.11.

With $m = 1$ only one sensor reading is used to determine the similarity between nodes. Hence, a WUR error will directly change the similarity information and with that the composition of the clusters. This will have an especially large impact when false wake-ups occur, as can be seen in Figure 4.11(b). Not only does the composition of the clusters change at moderate false wake-up rates ($\approx 8 \cdot 10^{-3}$), but also is it possible that a node is listed in several clusters, which can lead to a false data prediction. To mitigate this effect, a higher m is advisable if error rates greater 10^{-3} are expected. A good compromise between a short clustering period and a robust protocol is $m = 3$ and $thold = 2$.

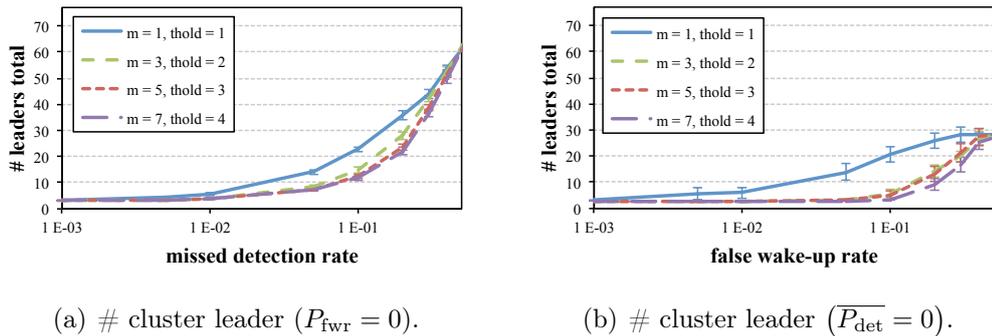


Figure 4.11: Total number of cluster leader for (a) the missed detection rate and for (b) the false wake-up rate.

4.4.2 Overall Protocol Simulation

In this subsection the whole protocol including the monitoring phase is analyzed. The first set of simulations evaluates the protocol when using the real measurements obtained from the ICE project [107]. The second set of simulations evaluates the behavior of the protocol when artificial sensor readings are used. With the artificial data model it becomes possible to evaluate the protocol for varying sensor readings, especially, to simulate a different data volatility to analyze the behavior of the protocol when a lot of reclusterings occur.

In the following the number of reclusterings is expressed in percent, where 100 % means that a reclustering process is initiated as often as possible—at the earliest three frames after the clustering phase is finished. Accordingly, a reclustering rate of 0 % means that after the clusters were established for the first time they were stable for the rest of the simulation—no reclustering was done. For the aeronautical use case a maximum loss of roughly 10 % of the data would be acceptable.

Using Real Measurements

The clustering protocol is evaluated when using real measurements from the ICE project [107]. Figure 4.12 shows the mean power consumption per node \bar{p} for all combinations of monitoring and clustering methods together with the corresponding goodput, but without the protocol extension. Additionally, mean power consumption and goodput is shown when not using clustering at all but just transmitting all measurements to the sink.

The average power consumption in the conventional cases is greater than when not clustering at all. And additionally, the goodput is slightly smaller than 100 % for the first and second monitoring method. For the third method the characteristics of the data were not matched. The number of reclusterings is very high (around 73 %). The loss of the data is not that high because only the knowledge about one packet per clustering is lost when using the conventional approach. But the average power consumption exceeds by more than five times the other monitoring methods. This clearly shows that the conventional clustering approach does not bring any advantage when using the set of real measurements. This is the case due to the high power consumption during the clustering phase.

When using a wake-up receiver for clustering, the lowest power consumption can be achieved when using the third monitoring method. Then, the nodes do not have to receive any packet and only have to transmit *outlier packets* if necessary. But, as can be seen, the characteristics of the data were not matched by the monitoring method. The number of reclusterings is very high which leads to an unacceptable loss of knowledge of about 50 % of the data. The difference between the first and second monitoring method is not

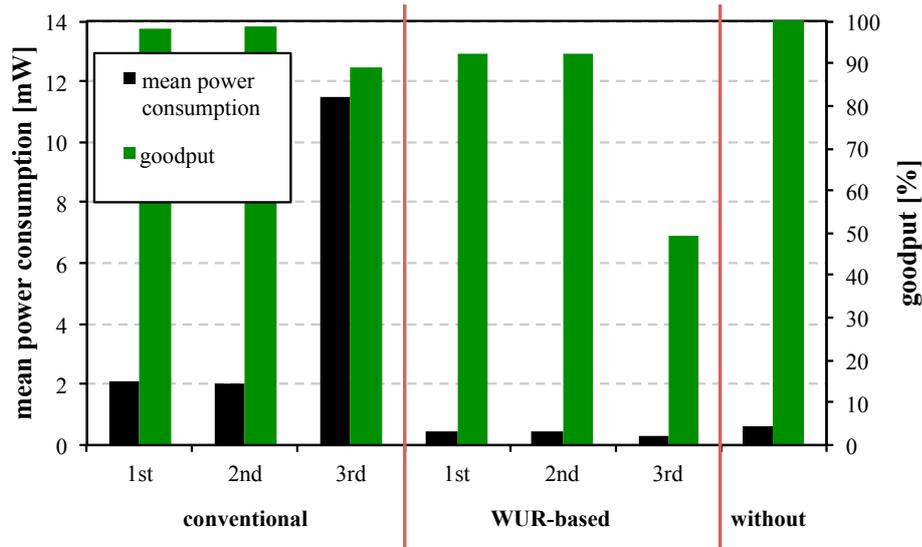


Figure 4.12: Mean power consumption per node and corresponding goodput when using the measurements of the ICE project [107]. Using the conventional as well as the WUR-based approach with all three monitoring methods.

distinguishable which indicates that the first method matches the characteristics of the data already well enough. The smoothing of the “cluster reading” has no negative influence on the goodput or the average power consumption. It only can have a positive effect in reducing the amount of necessary reclusterings. This positive effect, though, is bought with additional implementation complexity. The average power consumption when using no clustering at all is about 30 % higher than when using the WUR-based clustering protocol with monitoring method one or two. When using a CR2032 button cell battery (~ 200 mAh) as power supply this would extend the lifetime from 42 days to 55 days.

The extended protocol is also evaluated when using real measurements. But only for the WUR-based approach because the extended protocol will increase the power consumption and with an already too high power consumption a detailed analysis is not necessary. Using the first monitoring method results in a reclustering rate of 12.4 %, using the second method results in a reclustering rate of 12.1 %, and using the third method results in a rate of 81.1 %. As can be seen in Figure 4.13, the extension increases the goodput in all cases; in the case of the third method the increase is higher than 85 %. The average power consumption per node increases also when using the extension. Only slightly (8 %) for the first and second monitoring method but for the third method the increase is bigger (78 %). But the average power consumption for all methods is still below the benchmark—the

power consumption when not clustering at all. When using the extension the first two methods reach a goodput of 100 % and, therefore, outperform the benchmark. In this case, when using a CR2032 button cell battery, the lifetime would be extended from 42 days to 51 days while still keeping a goodput of 100 %.

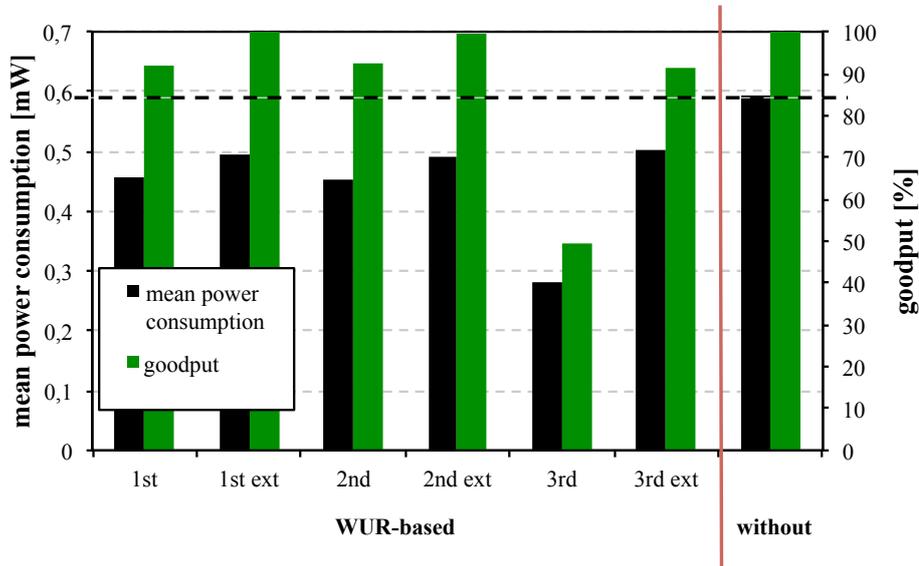


Figure 4.13: Average power consumption per node and corresponding goodput when using the measurements of the ICE project; including the protocol extension to transmit measurements generated during the clustering process.

Hence, using the real measurements from the ICE project, only a WUR-based clustering protocol would lower the energy consumption. Only the first two monitoring methods match the characteristics of the real measurements but when using the protocol extension even the third monitoring method provides an acceptable goodput of around 91 %. Thus, even if the monitoring method does not match the data well, when using the extension, the system can perform well.

Using the Artificial Data Model

Now the behavior of the protocol is analyzed when using the artificial data model. To simulate different numbers of reclusterings the state change probabilities of the Markov states are varied.

First, the average power \bar{p} per node is inspected for the conventional and for the WUR-based approach. The consumed power at a fixed reclusterings rate is nearly independent of the monitoring methods. Therefore, in Figure 4.14 only the graphs for the first monitoring method are shown. The goal of

different monitoring methods is to reduce the number of reclusterings by adapting the monitoring to the characteristics of the measurements. To get the same reclustering rate when using different monitoring methods different measurement characteristics have to be used.

The average power per node is calculated as the system-wide consumed energy divided by the total simulated time and the number of nodes. As can be seen in Figure 4.14, the number of reclusterings has a big influence on \bar{p} for the conventional clustering protocol, in which no wake-up receiver is used. But in the case of the novel clustering protocol the influence is only marginal. During the conventional clustering phase far more energy is consumed compared to when using wake-up receivers instead. With an increasing reclustering rate the energy-expensive conventional reclustering will have an increasing influence on \bar{p} .

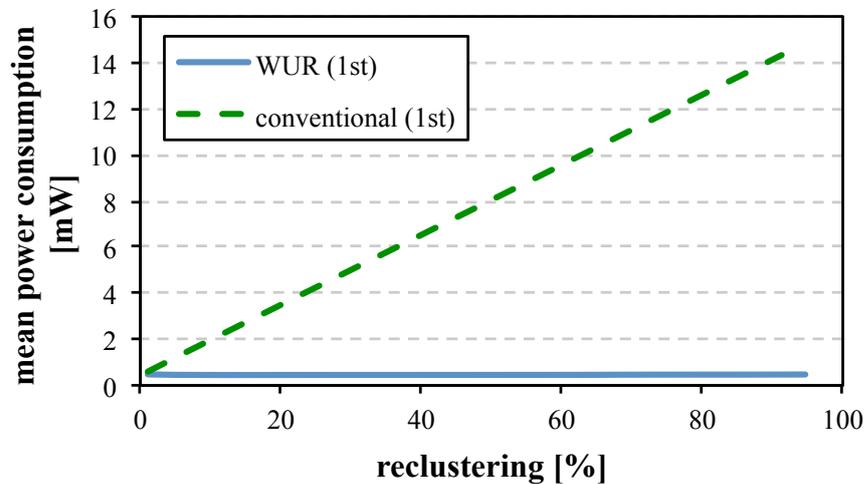


Figure 4.14: Average power consumption per node over reclustering in percent when using the artificial data model.

The consumed energy during clustering when using a wake-up receiver is that low that even 100% reclustering does not increase \bar{p} significantly. It is even lower than not clustering at all. A node has to spend more energy when just transmitting all sensor readings than it has to spend for reclustering. Therefore, in terms of energy consumption the wake-up receiver clustering protocol is free of charge.

In Figure 4.15 the resulting goodput is shown with and without using the protocol extension where measurement values generated during the clustering phase will be transmitted at a later point in time. When using a fixed data set, different monitoring methods just result in different reclustering rates, hence, no monitoring methods are shown; the goodput only depends on the reclustering rate.

Without using the extension, the goodput begins to drop instantly. For the conventional approach the maximum loss of knowledge is about 14 % and for the wake-up receiver approach the maximum loss is about 62 %. For the aeronautical use case a maximum loss of roughly 10 % would be acceptable.

When using the extension, in the wake-up receiver case the goodput is 100 % until a reclustering rate of approx. 60 % is reached. At this reclustering rate the average number of frames between clustering processes equals eight. A reclustering will be initiated when the sink received eight *reclustering requests* and a node will transmit a *reclustering requests* when three *outlier packets* were generated within five seconds. Hence, for at least eight nodes, three of these frames have to be occupied by *outlier packets* to initiate a reclustering. Therefore, only five frames are available to transmit the five measurements obtained during the clustering process. If the clustering rate increases further there are not enough slots available to transmit these measurements and the goodput will begin to drop.

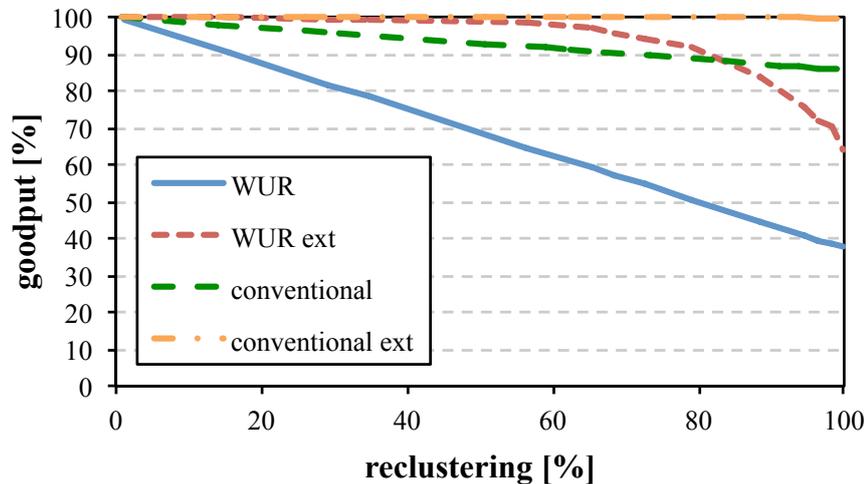


Figure 4.15: Goodput over reclustering in percent; with extension to transmit measurements obtained during the clustering process.

When using the extension with the conventional approach, only one measurement is not transmitted during the clustering phase. Therefore, the goodput will start to drop when the average number of frames between clustering processes equals four—which corresponds to a reclustering rate of approximately 90 %. Nevertheless, the goodput in the conventional case when using the protocol extension drops only to around 99.5 % because only one message per node has to be transmitted additionally.

Last, the average power consumptions per node is shown over different reclustering rates only for the WUR-based clustering approach. Again, because of the very high power consumption for the conventional approach

at even low reclustering rates, the conventional clustering approach is not depicted.

In Figure 4.16, the horizontal dashed black line represents the power consumption per node when not clustering at all but just transmitting all measurements. The graph for the exponentially smoothed monitoring method is not depicted in this figure, because it matches the graph of the first method. When using the extension the power consumption will increase faster with an increasing reclustering rate because more packets have to be transmitted. In the case of the third monitoring method the average power consumption is always lower than the benchmark line. In case of the first and second monitoring method the average power consumption in the extended mode is below the benchmark line until a reclustering rate of approximately 45%. At this reclustering rate the achieved goodput is still at 100%. The highest average power consumption occurs at a reclustering rate of approximately 80% and then decreases to a value slightly above the benchmark. At a reclustering rate of 80%, 10% of all nodes are barely capable of transmitting all their measurements obtained during the clustering process. After that point more nodes will not be able to transmit all of these packets because no further slots are available. Because of these dropped packets together with the low energy consumption during the clustering process itself the overall average power consumption will decrease.

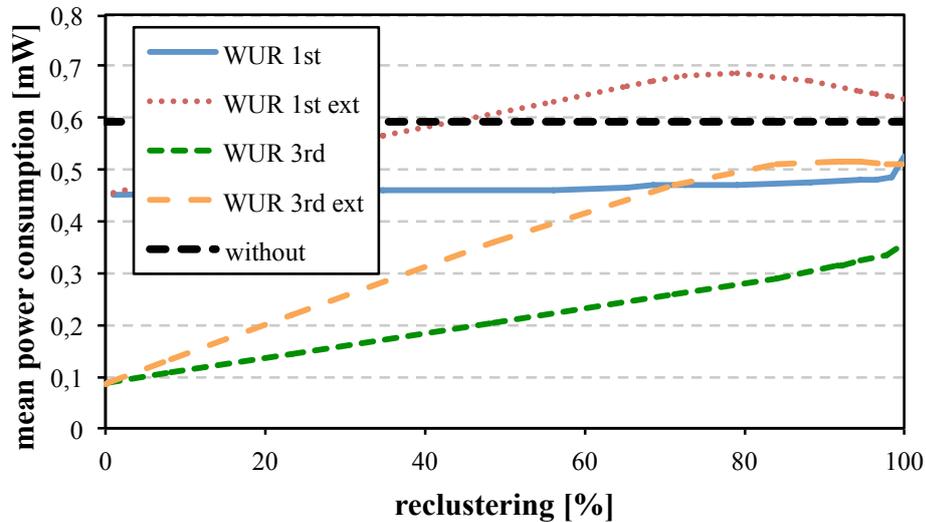


Figure 4.16: Average power consumption per node over reclustering in percent; with extension to transmit measurements obtained during the clustering process.

The characteristics of the measured data determine which monitoring method should be used. When the data is nearly constant for a long time the third

monitoring method would be the best solution owing to its low power consumption characteristics, but if not a high reclustering rate and, thus, a low goodput would result. If the data is volatile, depending on the required goodput, either the normal or the extended mode can be used with still moderate power consumption. The second monitoring method can result in a slightly lower reclustering rate at the cost of a higher implementation and computation effort. If goodput is not an issue the first two methods have the good characteristic of nearly constant power consumption at any given reclustering rate. If the goodput has to be high these monitoring methods still have lower power consumption until a reclustering rate of 45 %. At this reclustering rate, if a goodput of 100 % is required, the best solution would be to not cluster at all but just simply to transmit all measurements.

4.4.3 Verification by Hardware Implementation

Several wireless sensor nodes were assembled to test the clustering protocol in real life. The nodes consist of similar hardware devices as parametrized in the simulations and were equipped with either the wake-up receivers from the AETERNITAS project [39] or from the CHOSeN project [31].

Ten nodes were used to test the clustering protocol in real life. The presented clustering protocol is implemented as a task in FreeRTOS [109] and the temperature measurements from the ICE project [107] were saved on the sensor nodes to have a common data base for the purpose of comparison. These nodes were arranged in a circle with 2.6 m diameter. Additionally, another node was placed in the middle of it and was connected via USB to a PC to act as the data sink. The behavior of the WUR-based clustering protocol on hardware was compared to the simulated protocol behavior. In Figure 4.17 one clustering phase for both the simulated as well as the measured system is shown.

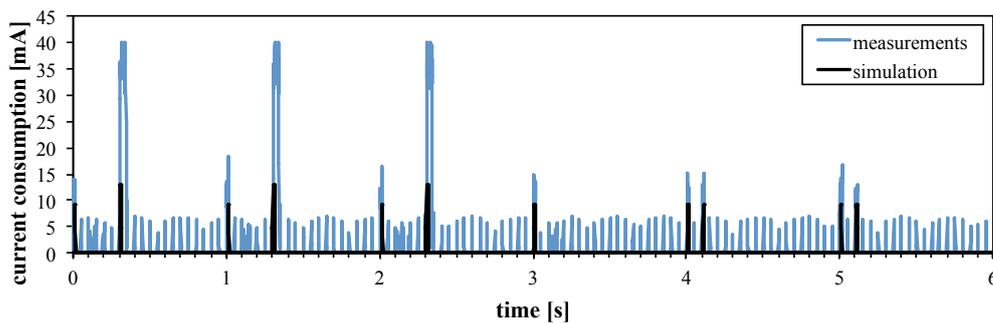


Figure 4.17: Comparison of the current consumption of real hardware measurements on a wireless sensor node and the simulated corresponding current consumption of the transceiver.

Two main differences are noticeable for the hardware measurements: the large peaks when transmitting the wake-up messages and the periodic small peaks at the beginning of each slot. The wake-up messages use on-off keying for modulation, which is not inherently supported by the AT86RF212 main transceiver. Therefore, the continuous wave mode of the transceiver was turned on and off by the microcontroller to modulate the wake-up messages. To reach the required modulation accuracy it was mandatory not to turn off the microcontroller. Hence, the required current consumption for transmitting a wake-up message was higher when using the hardware platform compared to the simulation. But, when using a transceiver which is capable of generating the wake-up messages by itself, the current consumption should be similar to the assumptions made in the simulation.

Additionally, the implementation of the TDMA protocol on the hardware nodes differed from the simulation. In the simulation the node was only woken up by the TDMA timer when an active slot was pending, while in the real implementation the TDMA timer woke up the node before each slot. With further optimizations the hardware TDMA timer can be adjusted to the same behavior but as a proof-of-concept this implementation suffices.

Besides those two differences, the simulation and the hardware measurements perform similarly. As can be seen in Figure 4.17, packets and wake-up messages were received and transmitted at the same points in time. Due to the increased current consumption for transmitting wake-up messages and due to non-optimized TDMA timers, the overall board power consumption was larger than in the simulation. But with a slightly changed board design and an optimized TDMA timer implementation the power consumption should be similar to the simulated one.

4.5 Conclusions

One of the prerequisites to use WSNs for flight applications is support for very dense sensor networks. It is highly likely that several nodes will measure correlated sensor readings, hence, the WUR-based clustering protocols were presented, which exploit this correlation. Simulation results show that the wake-up receiver enabled clustering and data transmission protocol outperforms conventional approaches (without wake-up receiver) in terms of energy consumption. This is valid for the real life data from aircraft cabin temperature measurements as well as for the artificial data model.

One main result is that clustering is free of charge (in terms of additional energy consumption) when using wake-up receivers and that energy consumption per node does not increase when the number of nodes in the network increases. The protocol is thus capable of increasing the life-time of networks with high node density and sufficiently high data transmission rates. If the sensor values from different nodes are correlated—as is the case of the inves-

tigated temperature measurements—the protocol saves approximately 75 % of energy per node while maintaining good data transmission reliability.

Variations in the measurement values greater than the cluster boundaries necessitate reclustering the network. When using the WUR-based clustering protocol without extension with monitoring methods 1 or 2, then the energy consumption does not increase with an increasing number of reclusterings. Additionally, even with a reclustering rate of up to 20 % the goodput remains reasonable. If reclustering rates of up to 80 % occur, the goodput can be kept at reasonable values when spending slightly more energy for the extended protocol mode. Hence, the WUR-based clustering protocol is more robust than conventional clustering approaches and outperforms the classical case of not clustering at all but just transmitting all sensor readings.

5 Increasing Transmission Reliability

High transmission reliability is one of the prerequisites for introducing WSN for flight applications. In this chapter a protocol is presented which introduces spatial diversity by using multiple access points per wireless cell. This protocol is evaluated in real world flight environments: in the upper stage of a space launcher Ariane 5, in an aircraft A330, and in an aircraft cabin mockup of an A330.

5.1 Introduction

As outlined in Section 2.1, many flight applications benefit from wireless sensor networks, e.g. Flight Test Instrumentation and Structural Health Monitoring. While the requirements are diverse in terms of data rate, delay, and loss tolerance, all applications should be served by the same wireless infrastructure and communication protocol. The protocol must be able to support many devices at a high density since there are potentially several thousands of wireless sensors in a single aircraft. For this purpose a robust MAC protocol for mission-critical applications is needed.

Hence, a protocol is presented that utilizes spatial diversity by using multiple access points per wireless cell. It is capable of increasing transmission reliability while increasing the resource demands only slightly.

5.2 Background

Utilizing spatial diversity by using multiple access points corresponds to macro diversity where several transceiver antennas are used for transferring the same signal. In macro diversity the distance between the antennas is much longer than the wavelength, hence, typically the antennas are located at different access points as opposed to micro diversity where the antennas are located at one device.

The received signals will then be combined to increase the probability of a successful reception. For example, macro diversity is utilized during a soft handover in Universal Mobile Telecommunications System (UMTS)—the third generation mobile cellular system. There, two received signals will be

combined when a mobile device establishes connections to two base stations. Macro diversity is also a research topic for wireless sensor networks [110–113]. For example, Valenti et al. [110] evaluate macro diversity used in a multi-hop relay channel scenario. Opposed to that scenario, in aircraft scenarios (see Section 2.2) a single-hop system is assumed, hence, macro diversity is only exploited for communication between nodes and access points; not for inter-node communication.

Several strategies are possible how to combine these received signals. Sulyman et al. [114] investigate the performance of generalized diversity selection combining (GSC) to conventional selection diversity combining (CSC) and maximal-ratio combining (MRC). In CSC just the strongest signal is chosen while in MRC all signals are combined with individual weightings based on their current SNR. MRC is the optimal combiner for independent additive white Gaussian noise channels but has a much higher implementation complexity.

To be able to use low-cost, low-power hardware devices implementation complexity has to be kept minimal. Hence, for this protocol selection combining is used as the diversity combining strategy. Furthermore, selection combining is used at packet level and not at bit level which results in simple selection of a successfully received packet (determined by cyclic redundancy check (CRC)).

5.3 Protocol Description

In dense WSNs with high traffic, contention is likely. Therefore, the MAC mechanism is a critical factor in the design of the system. TDMA has the advantage that channel capacity utilization, network delays, and power consumption can be determined in advance. The presented TDMA MAC protocol is designed to fulfill the requirements for dense WSNs for flight applications. This is achieved, in part, through the concept of AP redundancy; up to 8 redundant APs are supported per cell. The protocol uses a periodic time-slotted superframe structure that only consists of guaranteed time slots forming a contention-free period. A fixed pre-defined schedule will be used, hence, the number of APs has to be fixed before the schedule is calculated. Figure 5.1 shows such a TDMA superframe structure for this multi-AP protocol.

The superframe is divided into three parts: a beacon phase, a downlink phase, and an uplink phase. Beacons and downlink packets are sent once from each AP in the system while uplink packets are sent only once but are intended for all APs. When k APs are assigned to the communication cell the superframe consists of k beacons (one per AP), $m \cdot k$ identical downlink packets (k identical messages per AP for each downlink message m) and n uplink packets (one per node n). Therefore, the uplink exploits k times

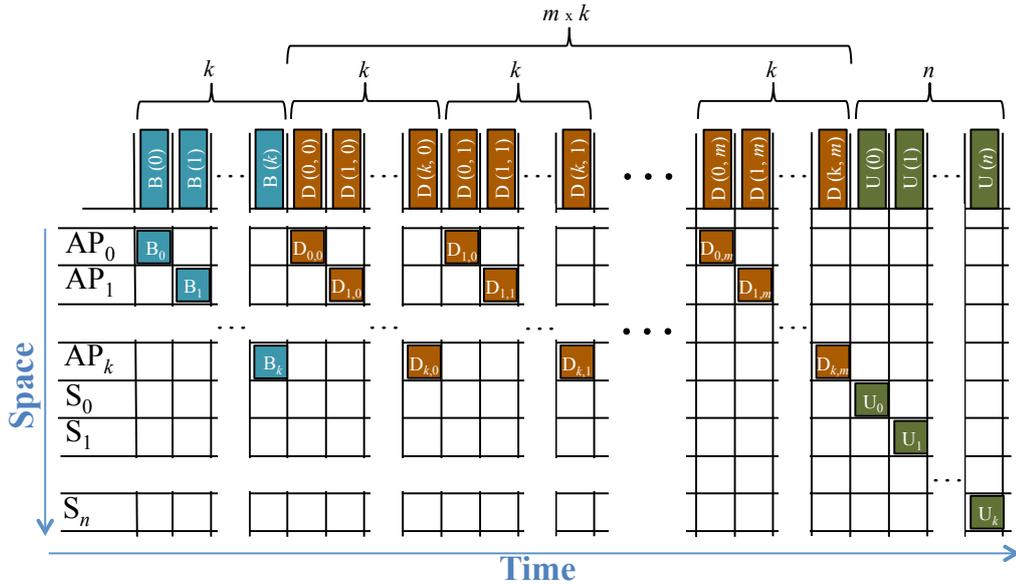


Figure 5.1: Multi-AP TDMA superframe structure depicted over time and space.

spatial diversity without using additional slots, while the downlink (beacons included) exploits k times spatial and k times temporal diversity, but uses k -times more slots to do so.

5.3.1 Network Synchronization

Both access points as well as nodes are synchronized by beacons. In each superframe each AP will broadcast a beacon which includes all necessary information to synchronize with the network. The AP with the lowest slot allocation—AP 0—will provide the initial synchronization to the rest of the devices. Thus, all other APs have to be in communication range of at least one lower-slot-allocation peer in order to synchronize correctly. By broadcasting k beacons per superframe, a reliable synchronization is most likely.

Before the network is established, all nodes will listen for an initial synchronization message. To save energy this is done using a very low duty cycle. In the very short ON-durations of this duty cycle they listen for power above a certain threshold on the channel, based on a clear channel assessment (CCA) algorithm implemented on the sensor nodes. Only when power above this threshold is detected, they stop duty cycling for $2 \cdot k$ slots to listen for beacons. If no beacon was detected they return to their duty cycle. To establish the network for the first time, all downlink and uplink slots are skipped for 1024 super frames, which consequently results in beacon bursts: $k \times 1024$ beacons will be broadcast. This rather long initial synchronization phase is used to increase the probability for a first time synchronization of

all sensor nodes.

When the network is established, it is sufficient for each sensor node to receive one beacon from any AP to stay synchronized. If a node does not receive any beacon within 10 superframes, it assumes itself as not synchronized anymore and stops sending uplink messages. Only after re-initializing the network such an unsynchronized node will be allowed to resynchronize and to join the network again.

5.4 Hardware Devices

Wireless sensor nodes and access points were specifically developed for aircraft applications. The AP was also developed as IP gateway and coordinator for the wireless protocol. Among its most important features, shown in Table 5.1, Power over Ethernet (PoE) facilitates the deployment of the network.

Table 5.1: Wireless access point data sheet

	Processor	ARM926EJ	
			ARM CM3 EFM32GG330F1024
	SRAM	4 MB + 128 KB	
	FLASH	8 MB + 1 MB	
	System	Embedded Linux + FreeRTOS	
	Transceiver	2.4 GHz AT86RF233	
	Power Source	PoE (Power over Ethernet)	
	Weight	47 g PCB	
	Dimensions		178 g incl. housing & antennas
			76 x 63 x 1 mm PCB
Temperature range		94 x 134 x 46 mm housing	
		-40 °C to +85 °C	

Similarly, wireless sensor nodes were specifically developed for flight applications. As already mentioned in Section 2.1, a very low power consumption is an important criteria, hence, to analyze the mean power consumption the duty cycle of the protocol is varied by modifying the superframe duration. With a longer superframe duration the duty cycle is reduced since sensor nodes are only active during beacon and downlink slots and during the transmission in their allocated uplink slot. The configured radio frequency (RF) parameters for the sensor node are -101 dBm receive (RX) sensitivity and $+4$ dBm transmit (TX) output power. The results of these measurements are presented in Figure 5.2.

With an increasing superframe length (lower duty cycle), the averaged current consumption approaches $22 \mu\text{A}$. This value is the minimum required current by the board when all components are in deep sleep mode. This

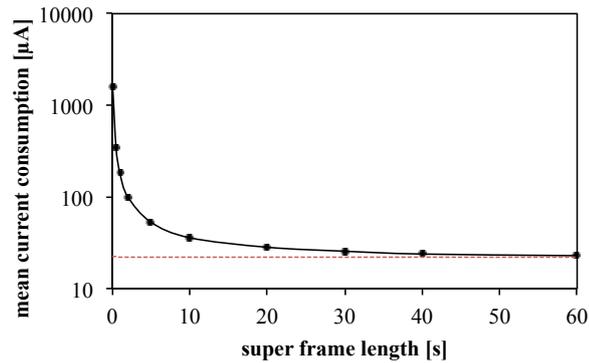


Figure 5.2: Mean current consumption of a wireless sensor node over a varying superframe length.

consumption is dominated by the power management chip necessary for the lithium polymer (LiPo) battery. The main features of the wireless sensor nodes are summarized in Table 5.2.

Table 5.2: Wireless sensor node data sheet

	Processor	ARM CM3 EFM32G210F128
	SRAM	16 KB
	Flash	128 KB
	System	FreeRTOS
	Transceiver	IEEE802.15.4 2.4 GHz AT86RF233
	Antenna	Fractus Compact Reach Xtend 1 dBi
	Built-in Sensors	Humidity and Temperature Sensor SHT21
	Power Source	Rechargeable Lithium Polymer 350 mA @ 3.7 V
	Weight	3.6 g PCB 21.4 g incl. battery & housing
	Dimensions	50 x 21 x 0.4 mm PCB 55 x 24 x 17 mm housing
Temperature range	-40 °C to +85 °C	

5.5 Measurement Campaigns

Experimental measurement campaigns have been performed with an implementation of this TDMA protocol. The aim was to characterize the wireless environment of flight applications and to confirm the improvements that the protocol provides.

5.5.1 Ariane 5

The first measurement campaign was undertaken in the connector torus of an upper stage of an Ariane 5 space launcher. The goal was to prove the functionality of the system and to establish an operational procedure for measurements in this great scale—six access points and around 90 wireless nodes were used within one wireless cell.

Setup

For the measurement campaign inside the upper stage of an Ariane 5 launcher the server and PoE switches were placed inside the hollow space of the launcher ring. The Ethernet cables to the access points were routed through holes in the inner shell of the launcher in order to connect the APs to the server. Six APs were distributed with approximately equal distance between neighbors inside the launcher torus. The APs were configured to use only one antenna. 70 wireless sensor nodes were distributed inside the ring in clockwise direction starting from AP 0. Additionally, 19 nodes were attached to the inner shell of the launcher torus. The nodes measure temperature, humidity, and battery voltage and transmit these values together with the measured Received Signal Strength Indicator (RSSI) values from the received beacons and downlink messages from each AP inside their allocated uplink slots.

As can be seen in Figure 5.3, the six APs are distributed in the launcher approximately equidistantly in the torus. Small variations are due to the installation environment, e.g. non-usable space where fuel tanks are located. 70 nodes are placed in clockwise direction starting at AP 0 with Node 1 and ending at AP 0 again with Node 70. Additional 19 nodes were placed at the outer wall of the torus in the middle of the launcher.

For the measurements the settings from Table 5.3 were used. After the initial synchronization phase in each frame all nodes send one uplink packet and each AP sends one downlink packet.

Measurement Results

The measured parameters are RSSI values for inter-AP communication, as well as for the downlink (AP to node) and for the uplink (node to AP). Additionally, in each frame all nodes measured the actual temperature, humidity, and battery voltage. First, the RSSI values for the inter-AP communication is inspected. The characteristics of the channels between each pair of APs is interesting insofar as they influence the synchronization of the network: AP 0 sends the master beacon, AP 1 synchronizes to that time signal, AP 2 needs to receive one of the beacons from either AP 0 or AP 1 and so on. Figure 5.4 shows the mean RSSI values from any AP to any other in both directions.

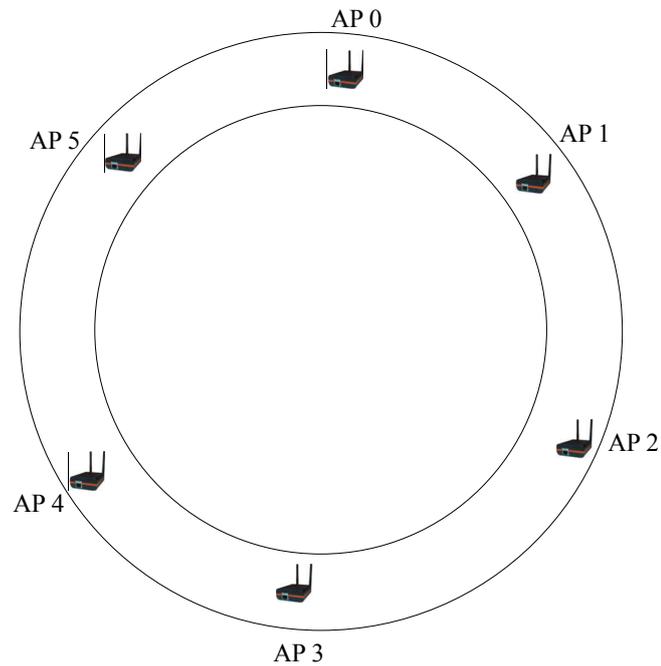


Figure 5.3: Placement of the six access points inside of the torus.

Table 5.3: Measurement setup Ariane 5

beacon slots	$k = 6$
downlink slots	$m = 1 \rightarrow 6$ slots
uplink slots	$n = 89$
slot length	10 ms
superframe length	1010 ms
physical data rate	250 kb/s
beacon size	27 byte
packet size	71 byte
transmit power	0 dBm
carrier frequency	2.4 GHz
physical layer	IEEE 802.15.4
channel	0x16

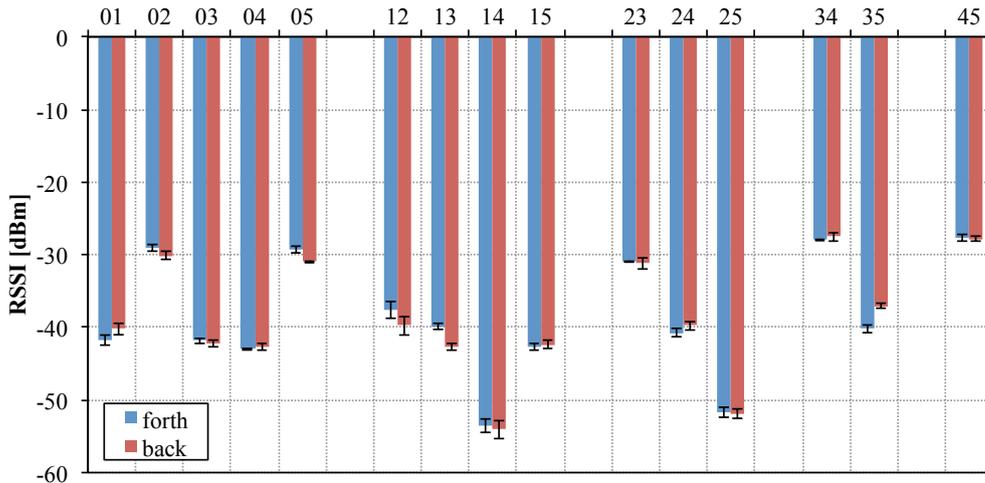


Figure 5.4: RSSI values between all combinations of access point pairs.

The error indicators at the bars show the confidence intervals of all RSSI measurements used for averaging. They are well within the accuracy/resolution limits of the hardware. The RSSI values are very high (always better than -60 dBm), from which follows that any AP can reach any other quite reliably if no further interference is present. The link between AP 0 and AP 1 is 10 dB worse than between AP 0 and AP 2, although the distance is shorter. This is due to the position of AP 0 with a metal tank in the line of sight to AP 1.

Signal strength of the nodes: Second, the signal strength between nodes and access points is examined. Figure 5.5 and Figure 5.6 show the averaged RSSI values for the uplink from all nodes to AP 0 to AP 5.

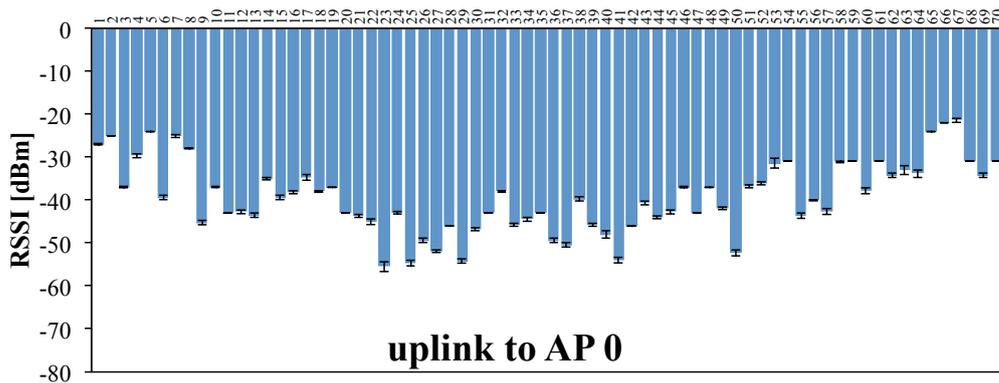


Figure 5.5: Signal strength from nodes to access point 0

The RSSI values correlate mainly with the distance except for fading and shadowing effects caused by obstacles in the wireless channel. An RSSI value

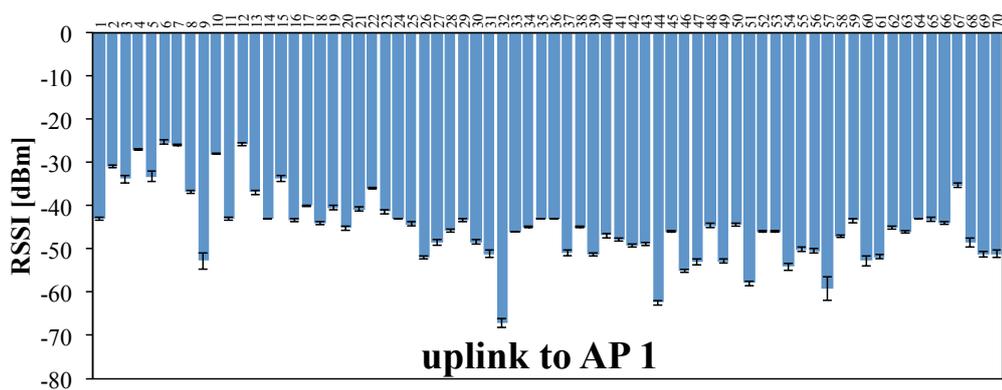


Figure 5.6: Signal strength from nodes to access point 1

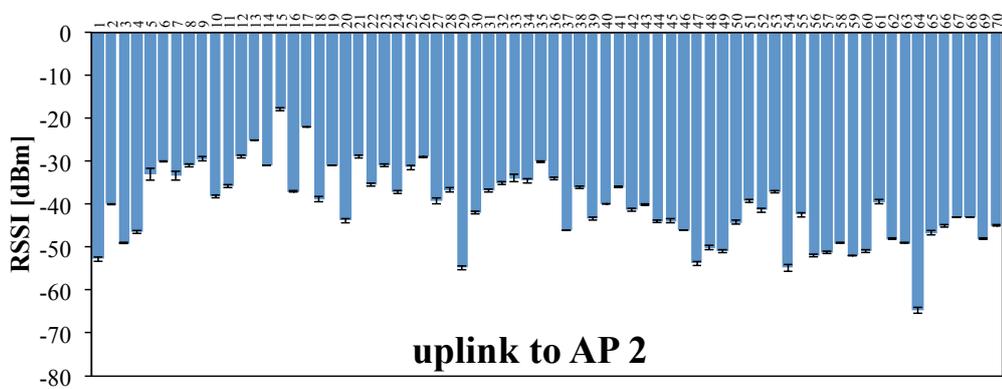


Figure 5.7: Signal strength from nodes to access point 2

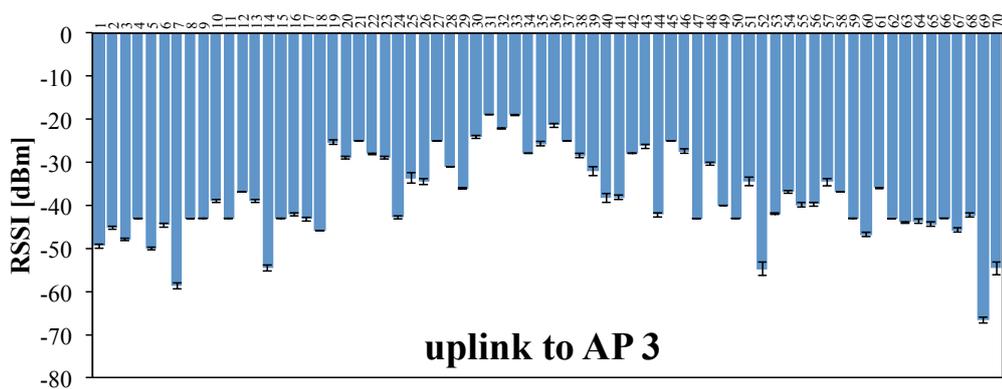


Figure 5.8: Signal strength from nodes to access point 3

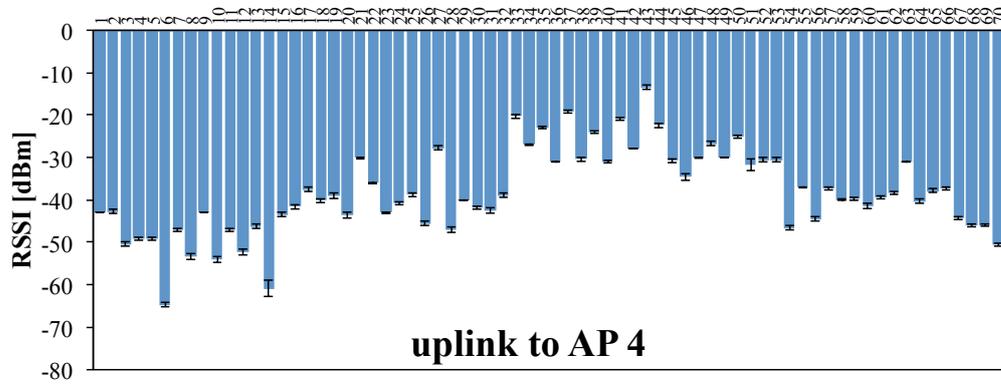


Figure 5.9: Signal strength from nodes to access point 4

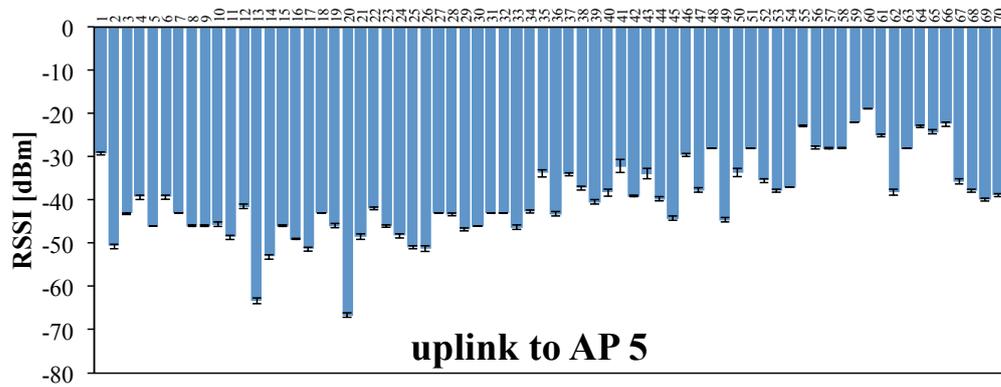


Figure 5.10: Signal strength from nodes to access point 5

better than -40 dBm indicates a very high communication quality, an RSSI value worse than -80 dBm leads to increased communication errors. Most of the nodes have a very performant link to all APs and all nodes have a very performant link to at least one AP. For example, Node 32 has an average RSSI of -67 dBm to AP 1 but -40 dBm to AP 0.

Considering only the link with the highest mean RSSI from a node to any of the APs, the average RSSI is always greater than -40 dBm, see Figure 5.11. Selecting the best RSSI from each node to only AP 0, 2, and 4 (Figure 5.12) leads to a very similar result, which means that 3 APs are sufficient to cover the launcher.

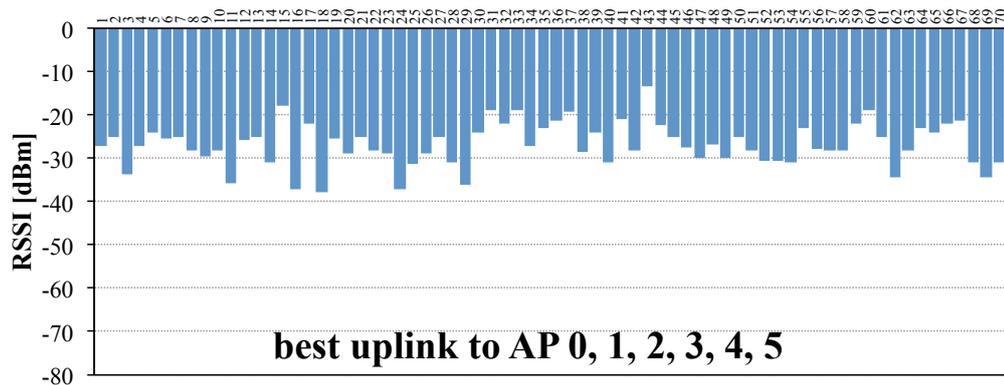


Figure 5.11: Highest signal strength from all nodes to one of the six access points.

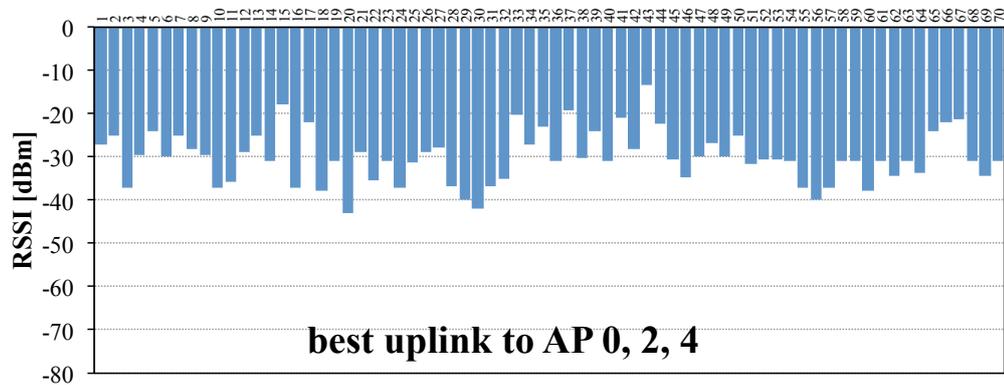


Figure 5.12: Highest signal strength from all nodes to one of the three access points 0, 2, or 4.

The standard deviation of the RSSI is very small (smaller than the accuracy of the measurement unit), which means that there is very little variation over time. That was to be expected due to the static environment. Variations

could have come from motion inside the launcher as well as from other wireless communication like wireless local area network (WLAN) or Bluetooth. These results already show that multiple redundant APs enable good communication quality for the overall system since any sensor can reach some or even all APs with very good signal strength. However, in order to quantify the reliability improvement coming from spatially redundant access points, the PER should be calculated. Unfortunately it was not possible to calculate the PER of the wireless system since, due to a software bug, there were also packets lost over the RS232 interface between the wireless module and the Ethernet module within the APs. These missing packets led to log files which cannot be evaluated in terms of PER. For a quantification of the PER, the measurement campaign has been repeated.

Conclusion Ariane 5 Measurements

The measurements show that one AP at any place in the launcher is able to reach most of the nodes with a signal strength of at least -60 dBm. However, due to non line-of-sight conditions the signal strength for some nodes is below -60 dBm. Hence, it is better to have redundant APs to be able to reach all nodes with this signal strength. Considering only the average RSSI values, three APs are a reasonable choice. The quantitative impact on the PER could not be calculated because of a serial interface bug in the APs. A follow-up measurement is necessary to measure the PER.

5.5.2 Airbus A330

During the Ariane 5 measurement campaign the functionality of the measurement system could be demonstrated and one aspect of the diversity concept was proved. But as was discovered a serial interface bug precluded the calculation of PERs. Hence, a second measurement campaign was undertaken in an aircraft Airbus A330 after fixing the serial interface.

Setup

The network was organized in three cells with two redundant APs per cell. The positions of the APs were selected, on the one hand, to provide maximum coverage and, on the other hand, to resemble likely installation positions. The sensor nodes were installed at the positions where the potential applications require them, for instance at the seats, life vest compartments, light units, lavatories, kitchens, galleys and cargo compartments. Each dot in Figure 5.13 represents one node; the colors symbolize different applications. However, during this measurements campaign, all of them are treated uniformly. The nodes were mostly placed on the left side of the aircraft due to the symmetry of its seat configuration along the aisle. Equipping all seats on the right side as

well would only lead to mirrored conditions regarding the opposed APs. The six red rhombi represent the APs. Two APs close together form a wireless cell and work on the same radio frequency. Assigning a different frequency for each cell enables interference-free, parallel operation of the cells.

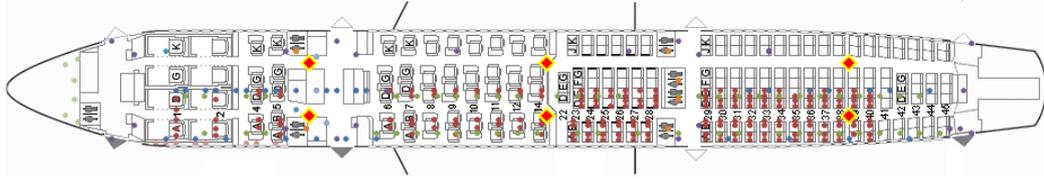


Figure 5.13: Setup of access points and wireless sensor nodes in the aircraft Airbus A330.

The sensor nodes are assigned to cells according to their distance from the APs with a maximum distance of 10 m. The aircraft is approximately 60 m long so that three cells with a radius of 10 m cover the whole aircraft. A maximum number of 170 nodes per cell was defined for the campaign, so that the superframe of a cell consists of two beacons (one from each AP), two downlink slots (again one from each AP) and 170 uplink slots (one for each sensor). With a slot length of 5 ms a superframe is 870 ms long. On average a packet is 33 bytes or 1 ms long which means that the duty cycle is $3/870 = 0.36\%$. In the worst case, when the first beacon is missed and also the command in the first downlink slot is missed, the node needs to listen for four whole slots leading to a duty cycle of $(4 \cdot 5 + 1)/870 = 2.4\%$. The measurement setup is summarized in Table 5.4.

Table 5.4: Measurement setup Airbus A330

beacon slots	2
downlink slots	2
uplink slots	170
slot length	5 ms
superframe length	870 ms
physical data rate	250 kb/s
beacon size	27 byte
packet size	40 byte
transmit power	0 dBm
carrier frequency	2.4 GHz
physical layer	IEEE802.15.4
channels	0x0B, 0x12, 0x1A

The measurement campaign was controlled by a central server. In Figure 5.14 the applicable periodic message sequence chart of the measurement protocol is depicted. The server triggers the APs via IP/UDP multicast

datagrams to broadcast the *measurement requests* according to the defined TDMA protocol. Nodes respond by sending their measured data in their slot and the APs forward the answers back to the server.

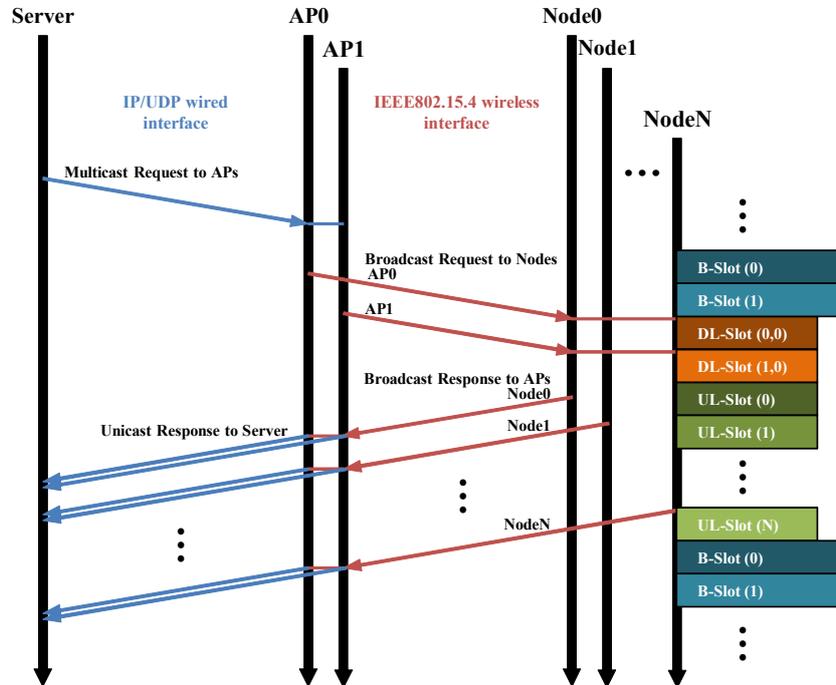


Figure 5.14: Message sequence chart for the Airbus A330 measurement campaign.

Each node measures temperature, battery voltage and RSSI values for the *measurement requests* received from each of the two access points. The APs measure the RSSI values for each received packet. Thus, four RSSI values are measured for each node per superframe ($AP_0 \rightarrow \text{Node}$, $AP_1 \rightarrow \text{Node}$, $\text{Node} \rightarrow AP_0$, $\text{Node} \rightarrow AP_1$). In total around 38 million transmissions were recorded. The packet error rates for the up- and downlink are calculated for each AP alone and for both combined. In the downlink, “combined” means a node received a request from either AP; and in the uplink, either AP received the packet from a node. Based on that the PER was calculated for the four links separately and for the combined links.

During the measurement campaign service technicians were installing and monitoring different aircraft systems inside the aircraft cabin. Hence, people were present during the measurements and were influencing the conditions of the signal paths. During a normal flight passengers and their luggage are present in an aircraft cabin as well, thus, service technicians and their equipment reflect the situation of a real flight.

Measurement Results

First, the quality of the wireless channel is evaluated with respect to RSSI values. Second, the correlation characteristic of the channels is evaluated and third the resulting PERs are investigated.

RSSI Values: The RSSI values obtained during the measurement campaign are in the range of -91 dBm to -10 dBm with a resolution of 1 dBm and an accuracy of 3 dBm, as stated in the transceiver data sheet [115]. In Figure 5.15, the average RSSI value per node over the distance between the access point and the corresponding node is depicted, together with four aircraft channel measurements from literature [9, 10, 12, 116] and the free-space model. As can be seen, the RSSI values form a dense cloud around the channel measurements from literature. But in contrast to them, a trend line for the RSSI values would follow an even steeper slope than suggested by the models. This is due to shadowing caused by real application positions, e.g. underneath the seats. Additionally, some nodes in places like the cargo compartment or the electronic bay have even lower RSSI values than the majority of the nodes. These nodes form the tight cloud below -70 dBm. Based on the measured distribution of the RSSI values we reason that a cell size of 10 m still leads to acceptable RSSI values. With a bigger cell radius the received signal strength will decrease to a level where it will have significant impact on the packet error rate for the current system.

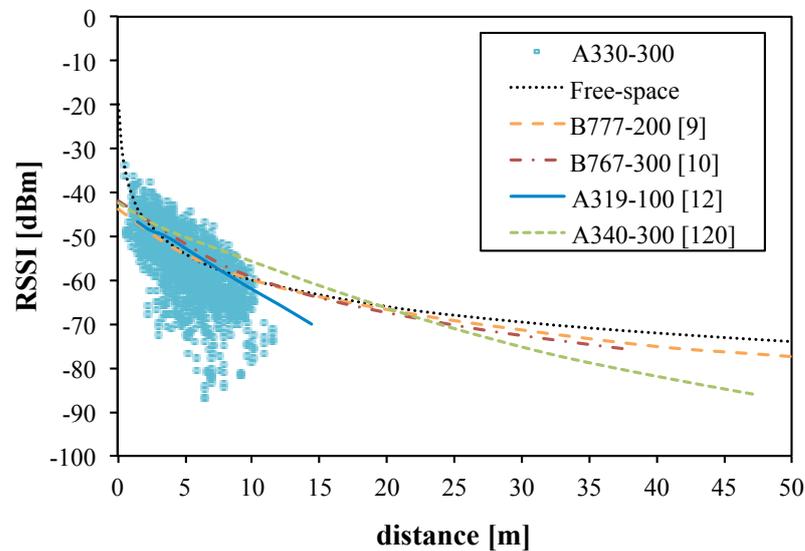


Figure 5.15: Average RSSI per node compared to channel measurements from literature [9, 10, 12, 116]

Figure 5.16 shows the difference of averaged RSSI values per node; in Figure 5.16(a) the difference of mean RSSI values between AP 1 and AP 0 and in

Figure 5.16(b) the difference of mean RSSI values between the downlink and the uplink are shown. The difference of the values can indicate if the channel is biased. In Figure 5.16(a), the RSSI values deviate around a common mean (≈ -2.1 dB); the received signal strength for AP 1 is slightly higher. As outlined in Section 5.5.2 the nodes are not equally distributed throughout the aircraft and therefore the distance to Access Point 1 is a little bit shorter on average; hence the higher RSSI values. Nevertheless, there are lots of nodes for which the signal strength is higher at AP 0. This emphasizes the benefit of the multi-AP approach. In Figure 5.16(b) the difference of the RSSI values of uplink – downlink is shown. The deviation of the difference is smaller than in Figure 5.16(a) and in the region of the accuracy of the RSSI values. We expected this deviation, but with a mean of 0 instead of -2.6 dB. The RSSI values for the downlink were always higher than for the uplink. Most probably it is caused by non-reciprocal components in the access points, but has to be validated with further measurements.

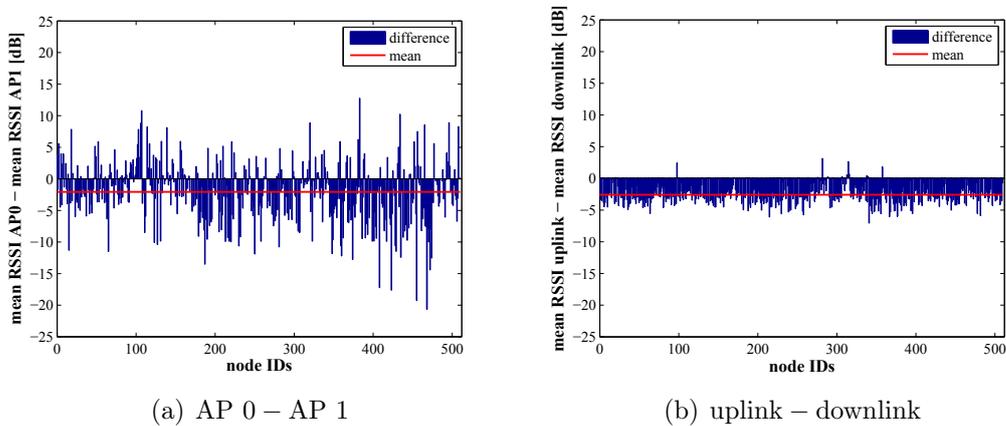


Figure 5.16: Difference of the mean RSSI values of (a) AP 0 and 1, and (b) uplink and downlink.

Channel Correlation: Two access points are responsible for each cell. Downlink packets—beacons included—are sent by both access points in different slots. Uplink packets are broadcast and because of the cell size can be received by both access points. Therefore, in both directions spatial diversity and in the downlink also temporal diversity is exploited. To gain as much diversity as possible, the best case would be when the links to AP 0 and AP 1 were statistically independent. Then, the combined PER would equal the product of the individual PERs of the APs, ($PER_C = PER_{AP0} \cdot PER_{AP1}$). To evaluate the dependence of those links, we calculated the cross-correlation of the RSSI values and determined the *Pearson product-moment correlation coefficient* ρ per node. A high correlation coefficient means that both APs

measure similar RSSI values, while a coefficient of zero implies that there is no linear correlation between the measurements. Figure 5.17 visualizes this dependency. The links for the depicted node in Figure 5.17(a) are statistically dependent; both APs measure similar RSSI values at each time, which therefore form a cloud around the bisecting line. Hence, a high correlation coefficient of $\rho_{0,1} = 0.90$ can be observed for the values of AP 0 and AP 1. Nearly statistically independent links are depicted in Figure 5.17(b). For this node, the RSSI values do not follow the bisecting line but form a loose cloud ($\rho_{0,1} = 0.04$).

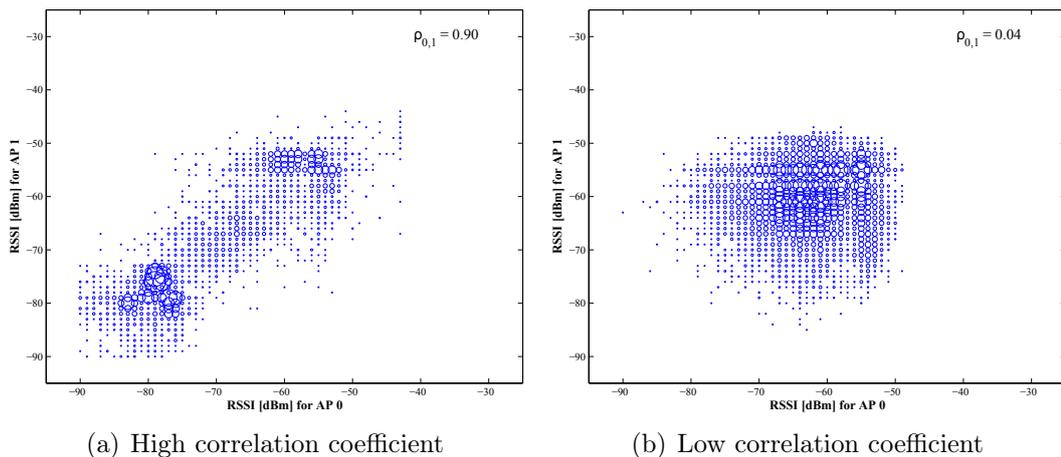


Figure 5.17: RSSI values for AP 0 and 1 in the downlink; (a): RSSI values for Node 165, and (b): RSSI values for Node 45. The diameter of the circles is proportional to the number of measurements.

For each node such a correlation figure was generated and categorized. Three categories were observed. The first category contains all the cases where the channel conditions for both APs are stable. Or more precisely were over 90% of the RSSI values are in a range of 10 dBm. This category is the least frequently occurring category which contains roughly 18% of all nodes. As an example for this category, Node 485 is depicted in Figure 5.18.

When the channel is stable for both APs it would be enough to use just the one AP which has the better channel conditions. The second access point does not increase the transmission reliability significantly. Nevertheless, it might not be known beforehand which access point will have the better channel conditions, hence, this access point redundancy increases transmission reliability.

The second most occurring category with around 33% contains all the cases when only one access point has a stable channel condition and the second one has an unstable condition. Such a case is depicted in Figure 5.19: in (a) the channel condition of AP 0 is stable for Node 410, and in (b) the channel condition for AP 1 is stable for Node 16.

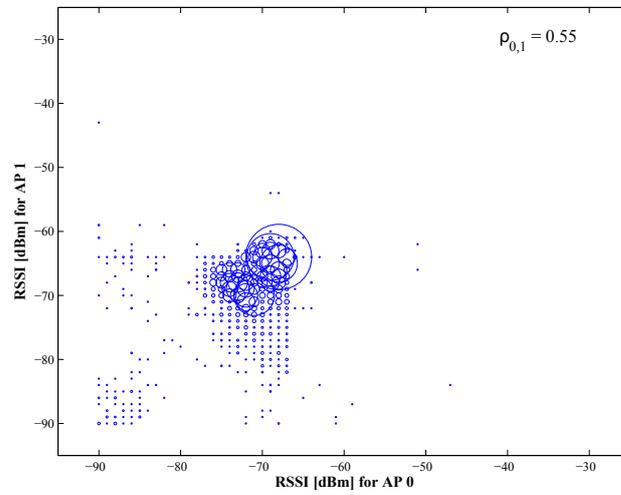


Figure 5.18: Category I: stable channel conditions for both APs (Node 485 is depicted).

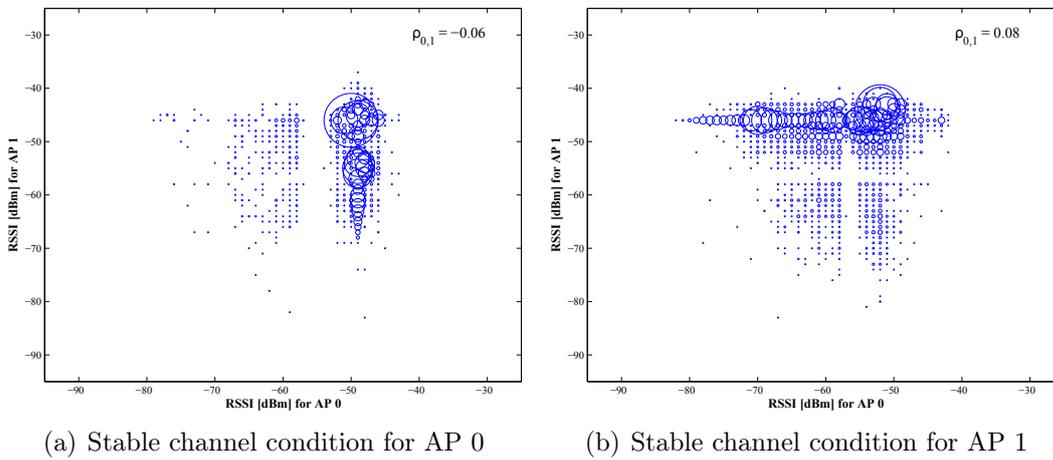


Figure 5.19: Category II: stable channel conditions for only one AP; either for (a) AP 0 (Node 410), or (b) AP 1 (Node 16).

Similar to Category I, if the access point with a stable channel condition has a good characteristic and is known beforehand, the second AP would not be necessary. Normally, this is not the case. Additionally, if the stable channel has only a medium channel characteristic, it can happen that the second AP has a better condition sometimes. Hence, the diversity helps, as well.

Last, the most frequently occurring category with around 45% contains all the cases when most of the time both access points do not have extreme channel conditions at the same time. In these cases the RSSI values form a cross in the correlation figures with most of the values at the middle of the cross. Such an instance can be seen in Figure 5.20. Such a cross means that the probability that both access points experience an extreme channel condition at the same time is very low; neither very good nor very bad channel conditions are experienced at the same time. This behavior is caused, most probably, by the service technicians inside the aircraft. Someone will stand in the LoS of one channel but the probability is high that the LoS for the second channel is not obstructed. Hence, this cross-behavior occurs. Here, clearly the redundant AP helps to increase the transmission reliability.

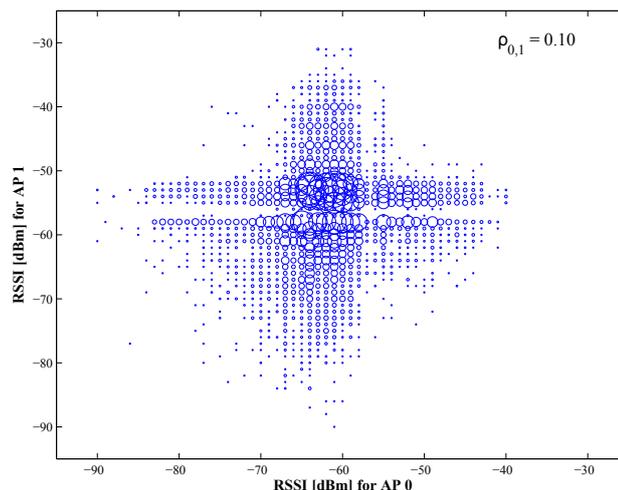


Figure 5.20: Category III: most of the time, both access points have different channel conditions (depicted for Node 332).

In general, the links to both APs are slightly correlated, but this correlation is still low enough to exploit diversity. The mean correlation coefficient over all nodes in the downlink equals 0.30 with a standard deviation of 0.42, while for the uplink the mean equals 0.21 with a standard deviation of 0.31. The high deviation indicates a different potential gain for each node by the multi-AP approach; nevertheless, for most of the nodes the impact will be high.

Packet Error Rates: The RSSI values are only important as a method to estimate the channel conditions and the resulting PERs. Hence, in the

following the PERs itself are investigated; for each link separately as well as the combined PERs.

Figure 5.21 shows the packet error rate for all nodes in the uplink case. There, the nodes are sorted by their PER values in decreasing order for each AP separately. This means that a vertical line does not compare the PER for one specific node, but rather shows the scale of possible PERs per AP. As can be seen, for most of the nodes the PER is below 10^{-2} with the exception of some outlier nodes. The moderate outliers occur when the link quality is poor due to “challenging” placements of the nodes. The outliers with high values result from nodes which lost connectivity to the APs for a long time. For all other nodes the PERs vary from 10^{-2} to 10^{-4} . But in the combined case, where a packet is only lost if no AP was able to receive it, the PERs could be reduced by roughly one order of magnitude.

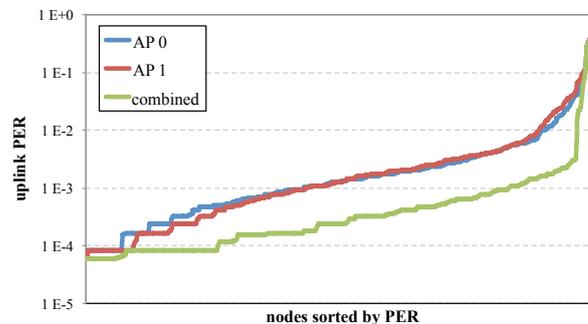


Figure 5.21: PER for all nodes sorted by their PER value for each AP individually for the uplink.

Figure 5.22 shows the median of the PERs for the downlink, the uplink, and for the round-trip communication. The PER for AP 1 is in all three cases slightly higher than the PER for AP 0. In the downlink 18%, in the uplink 9%, and for the round-trip 17%, even though AP 1 received a 2.1 dB stronger signal. As can be seen in Figure 5.15, the RSSI values are in the region of -40 dBm to -70 dBm, therefore, even with a 2.1 dB weaker signal the signal-to-noise ratio is still high and will have no big influence on the PER. The same is valid when comparing the downlink with the uplink. The 2.6 dB stronger signal in the downlink also has no significant influence on the PER.

The combined PER will be lowest if both links are statistically independent. With the given PERs for AP 0 and AP 1, this lower bound is on the order of 10^{-6} . The measured combined PER was on the order of 10^{-4} . Nevertheless, in all three cases the combined PER is much lower than with only one AP, which clearly shows the benefit of the multi-AP approach when introducing spatial diversity. The round-trip PER is of course higher than only for the down- or uplink. Nevertheless, the combined PER in the round-

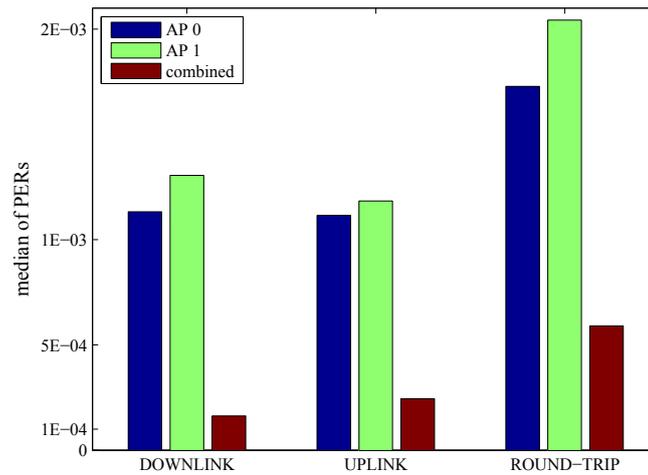


Figure 5.22: Median of PER for AP 0, 1, and combined.

trip case is still lower than the PERs for AP 0 or 1 in down- and uplink individually.

Conclusions Airbus A330 Measurements

A wireless sensor network with 500 sensor nodes was successfully deployed in an Airbus A330-300. Robustness is achieved by spatial diversity provided by multiple APs. The measured environment is even more challenging than suggested by channel models from literature, which becomes obvious when comparing the measured distribution of RSSI values with the models. The reason for that is the deployment of the nodes at the application positions, e.g. in the life vest compartments under the seats. Two access points per cell were used to introduce spatial diversity. This diversity reduced the PER significantly.

5.5.3 Airbus A330 Mock-up

The last measurement campaign was undertaken in an Airbus A330 Mock-up. The great advantage of an aircraft mockup over a regular aircraft was the possibility to perform long-term measurements. The goal of this measurement campaign was, on the one hand, to further characterize the aircraft channel with regard to PER performances when using several APs and, on the other hand, to establish a packet error rate channel model which includes distributions of burst error lengths and run lengths.

Setup

The measurements were performed in an Airbus A330 aircraft cabin mockup. In this mockup, six APs and 52 nodes were arranged as depicted in Figure 5.23. Two nodes were placed at each life vest compartment under each seat, one node was placed at each arm-rest and one node was placed at each passenger service unit directly above each seat. Four APs were placed in the hatrack and two were placed on floor-level in the middle row. The IDs of the APs were assigned in the following order: as visualized in Figure 5.23(a), AP 0 is in the top left, AP 1 in the top right, AP 2 in the bottom left, AP 3 in the bottom right, AP 4 in the center top, and AP 5 in the center bottom. These six APs form one wireless cell inside the mockup and are in communication range of each other. The maximum distance between them is around 6.6 m. For all measurements the same devices as described in [7] were used.

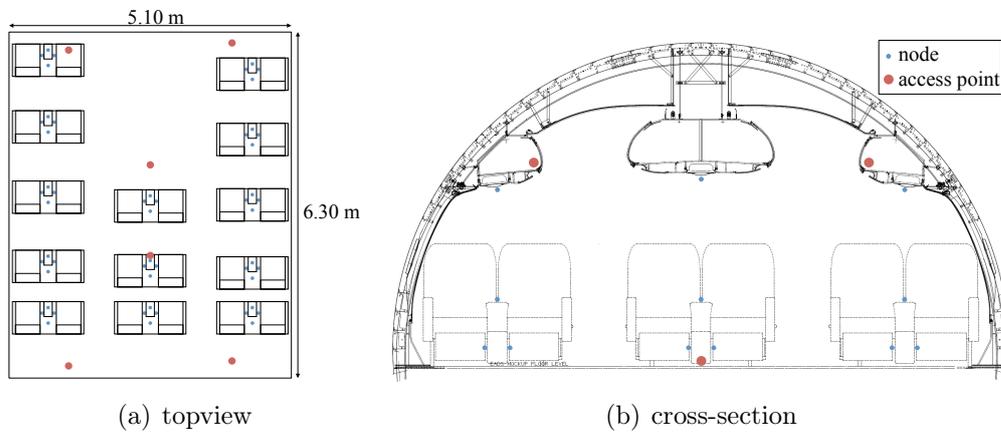


Figure 5.23: Airbus A330 Mockup with placement of nodes and APs.

For the TDMA protocol a setup as depicted in Table 5.5 was used. The nodes with the IDs $\{0,1,12,18,30\}$ were not used, hence, no packets were sent in their slots. Two measurement runs were done: the first run took over three days without people inside the mockup to characterize an empty environment. In the second run a typical dynamic environment was characterized. For two hours people were inside the cabin to simulate a flight (boarding and de-boarding included). Depending on the measurement length different resolutions for the packet error rates can be obtained. In the empty case around four hundred thousand packets per node were transmitted resulting in a PER resolution of $2.69 \cdot 10^{-6}$, in the manned case over five thousand packets per node were sent which results in a much lower PER resolution of $1.44 \cdot 10^{-4}$.

Several measurements are performed in each superframe. Each node measures the temperature and battery voltage as well as the RSSI for all received downlink packets from the APs. These measurements form the payload of the uplink packets. Accordingly, the APs measure the RSSI of each received

Table 5.5: Measurement setup

beacon slots	6
downlink slots	6
uplink slots	58
slot length	10 ms
superframe length	700 ms
physical data rate	250 kb/s
beacon size	27 byte
packet size	71 byte
transmit power	0 dBm
carrier frequency	2.4 GHz
physical layer	IEEE 802.15.4
radio	AT86RF233 [115]

uplink packet. In total 12 RSSI values are measured for each node per superframe. Additionally, all packets received by the APs are stored at a central data logger in order to evaluate the packet error and bit error rate characteristics.

Measurement Results

First, the difference of the wireless channel in the empty and manned case is highlighted with respect to its RSSI characteristics. Second, the impact of the multi-AP concept is evaluated regarding combined packet error rates. Last, bit error rates are investigated to evaluate the possibility of using additional error correcting codes.

RSSI Values: As stated in Section 5.5.3 the RSSI values for both the downlink and the uplink direction were measured. The difference between the uplink and downlink was only marginal for each measurement. But the difference in the RSSI characteristic between the empty and the manned aircraft cabin was obvious. Hence, for this measurement campaign, only the uplink characteristics are analyzed, but for both the empty and the manned scenario. In Figure 5.24, as an example, the uplink RSSI values between Node 5 and AP 4 for both scenarios are depicted. As can be seen, in the empty case the RSSI values are very stable and only fluctuate around the accuracy of the RSSI values of 3 dB, while in the manned case the RSSI values fluctuate very strongly in the range of about 15 dB. These strong fluctuations in the manned case originate from the movement of the passengers within the cabin; sometimes the direct line of sight will be blocked by a passenger, hence, the decline in the received signal strength.

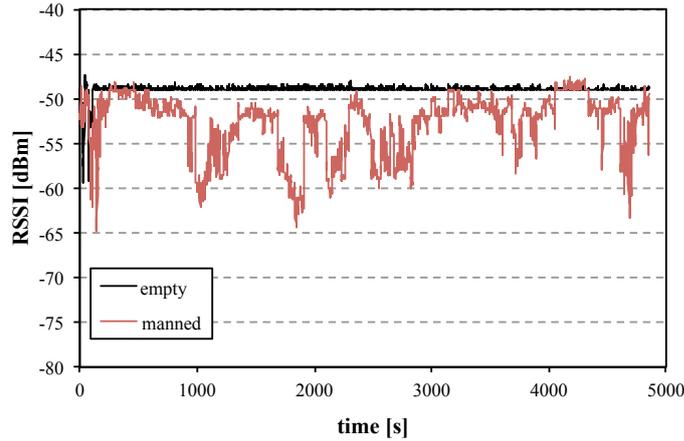


Figure 5.24: Uplink RSSI between AP 4 and Node 5 for the empty and manned scenario.

Packet Error Rates: From this decline in signal strength, an increase of the PER does not follow necessarily; the PER has to be inspected individually. In Figure 5.25 the packet error rates are shown for all individual AP links. There, the nodes are sorted by their PER values in decreasing order for each AP separately. This means that a vertical line does not compare the PER for one specific node, but rather shows the scale of possible PERs per AP. In Figure 5.25(a)—the empty scenario—the PER varies from $1 \cdot 10^{-6}$ to $1 \cdot 10^{-3}$, where a PER of $1 \cdot 10^{-6}$ corresponds to the possible PER resolution. Each AP reaches some nodes with a low PER and others with a high PER, but for most of the nodes the PER resolution can be reached. For each individual AP the distribution of the PER is similar.

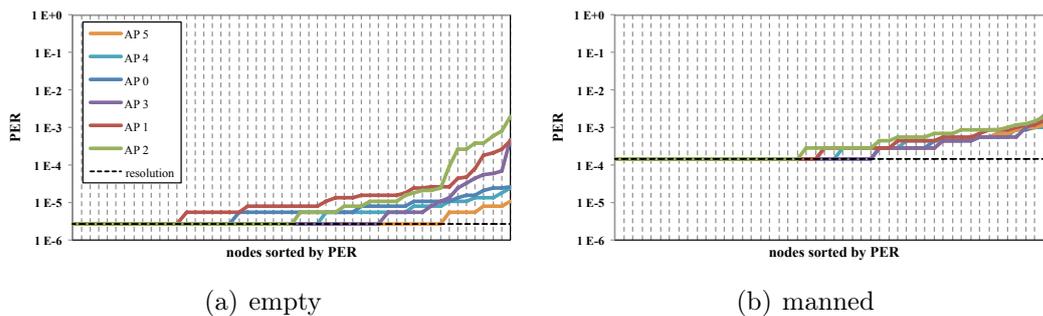


Figure 5.25: PER for all nodes sorted by their PER value for each AP individually. Both, for the (a) empty and (b) manned scenario.

In Figure 5.25(b) the distribution of the PERs per node is shown for the manned scenario. As can be seen, the PER varies in the same region as for the empty scenario (within its own resolution of course). But more nodes

than in the empty scenario have a worsened PER. Roughly half of the nodes per AP still reach the PER resolution, which indicates that the people in the mockup do not influence the channels to all nodes. But the nodes which are influenced have a clearly increased PER.

In Figure 5.26 the combined PER per node is shown. As an example these PERs are investigated for some AP group combinations. The chosen groups just show the distribution of the PER over the combinations, a more complete comparison is done later on. As can be seen in Figure 5.26(a), when using more than one AP the individual PER per node can be decreased for most of the nodes. For instance, when inspecting the PER for the AP group combination $\{\text{AP 0, AP 1, AP 2}\}$ it can be seen that for all nodes the PER is decreased to at least $8.06 \cdot 10^{-6}$ and for around 78% of the nodes the PER resolution can be met. The overall PER was clearly improved. This becomes even more clear when inspecting the manned scenario. There, most nodes suffer from a higher PER due to the dynamic channels. But when using more than one AP the PER decreases for all nodes. When using three APs for almost all nodes it was possible to reach the PER resolution.

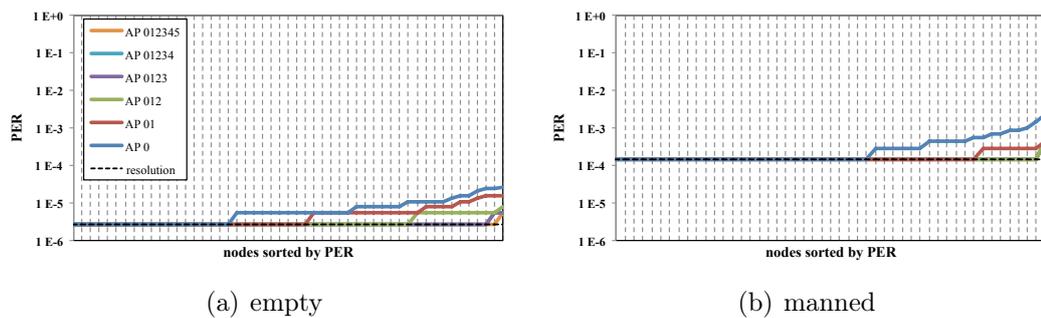


Figure 5.26: PER for all nodes sorted by their PER for each AP group individually. For both, (a) the empty and (b) the manned scenario.

To inspect the general behavior when using more than one AP the combined PER is now evaluated. Therefore, the combined PERs of all individual nodes to each group of APs are averaged. The standard deviation of these mean values is very high. Hence, the mean PER is more a measure for the general performance of an AP group than the representation of the individual PER per node. These mean PERs for all possible AP group combinations are shown in Figure 5.27.

The AP combinations are divided into six blocks with an increasing number of APs per group. Within a block all combinations with linearly increasing access point IDs are shown. For example, for Block 2 it starts with the AP combinations $\{\text{AP 0, AP 1}\}$, $\{\text{AP 0, AP 2}\}$, $\{\text{AP 0, AP 3}\}$, etc. Also depicted is the mean PER per block, which is the second order mean—the mean of the AP group values. As mentioned before, because of the broad range of PERs

per AP, the standard deviation is very high and in the order of the mean PER. Nevertheless, the mean PER per block—PER(x)—gives a good estimation for the performance of the system. As can be seen in Figure 5.27(a), in the empty case the mean PER can be reduced from $2.77 \cdot 10^{-5}$ when using only one AP to the PER resolution when using three APs. In the empty scenario more than three APs do not improve the system further. In Figure 5.27(b) the manned scenario is evaluated. The same behavior can be seen as in the empty scenario. The poor PER of $3.94 \cdot 10^{-4}$ when using only one AP can be improved to the PER resolution when using three APs. Again, more than three APs do not improve the PER further. But with a higher PER resolution (a longer measurement) the difference between the group PERs could be inspected in more detail. This clearly shows that with spatial diversity the high volatility of the channels in the manned scenario can be compensated for and the same region of PER can be achieved as in the empty scenario.

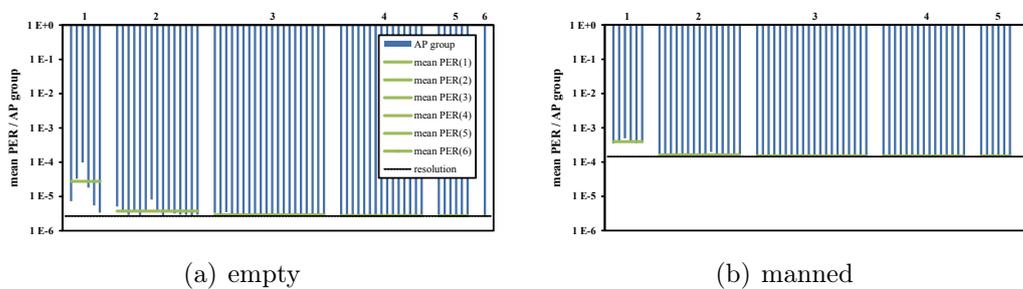


Figure 5.27: Mean PER for all AP group combinations. For both, (a) the empty and (b) the manned scenario.

Packet Error Distributions: To create a statistical model for the packet error rate in the aeronautical channels the burst error length and the run length for both scenarios were evaluated. With these two statistics the channel characteristic can be determined. The burst error length is the number of consecutive erroneous packets and is a measure for determining the duration of channel errors. In Figure 5.28 the relative frequency of the burst error length is shown. As can be seen, for both cases a burst error length of one has the highest probability for all APs; nearly 100%. Additionally, the maximum occurred burst error length was three. In the empty scenario only for AP 2 with a probability of about 0.1% and in the manned scenario for AP 3 with around 1%. The probability for a burst error length greater than one is for all six APs very small. For instance, in the empty scenario only in the case of AP 2 and AP 3 a burst error length longer than one occurred. In the manned case, a burst error length longer than one occurred for all APs, but for most with a probability around 1%. This clearly underlines that only short error bursts occur. In the manned scenario the probability

for greater burst error lengths is slightly higher than in the empty scenario, but is following the same behavior.

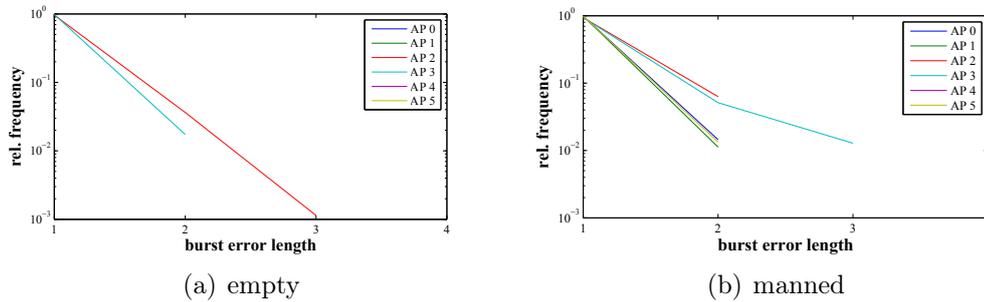


Figure 5.28: Relative frequency of the burst error length for both, (a) the empty and (b) the manned scenario.

The opposite channel characteristic is evaluated in Figure 5.29. There, the complementary cumulative distribution function (CCDF) of the run length is depicted. The CCDF is the inverse of the CDF and describes the probability for a run length greater than x : $F(x) = P(X > x)$. As can be seen, the probability for a run length greater than 10,000 is around 50% in the empty scenario and about 30% in the manned one. Additionally, in the manned case the run lengths are shorter. There, the longest run was around 10,000 packets while in the empty scenario it was around one million packets without error. With both the burst error length and run length characteristics, a Markov chain with corresponding transition probabilities can be used to model the channel behavior.

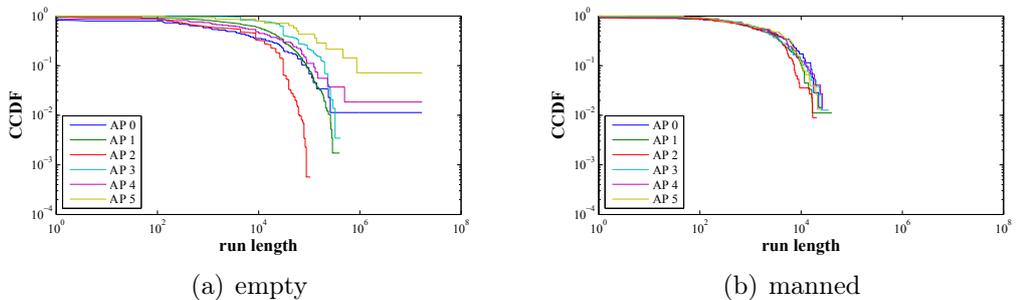


Figure 5.29: Complementary cumulative distribution function (CCDF) of the run length for both, (a) the empty and (b) the manned scenario.

As shown, the multi-AP concept reduces packet errors which occurred due to dynamic channel behavior, for instance, caused by people which are blocking the direct line of sight between one AP and a node. The multi-AP concept does not decrease the resulting PER to a relevant degree if the errors are correlated for all APs. Then, another possibility to further decrease the

PER would be to use additional error correcting codes. Therefore, the bit error characteristics are evaluated.

Bit Error Characteristic: In Figure 5.30 the absolute frequency of bit errors per erroneous packets is depicted. The most frequent number of bit errors is about 10 with decreasing probability around it. Additionally, in Figure 5.31 the cumulative distribution function (CDF) of bit errors is shown. The packets to some APs experience more bit errors than others. Also, 50% of the erroneous packets have between 1 and 60 bit errors and for around 90% of the packets not more than 100 bit errors per packet occur.

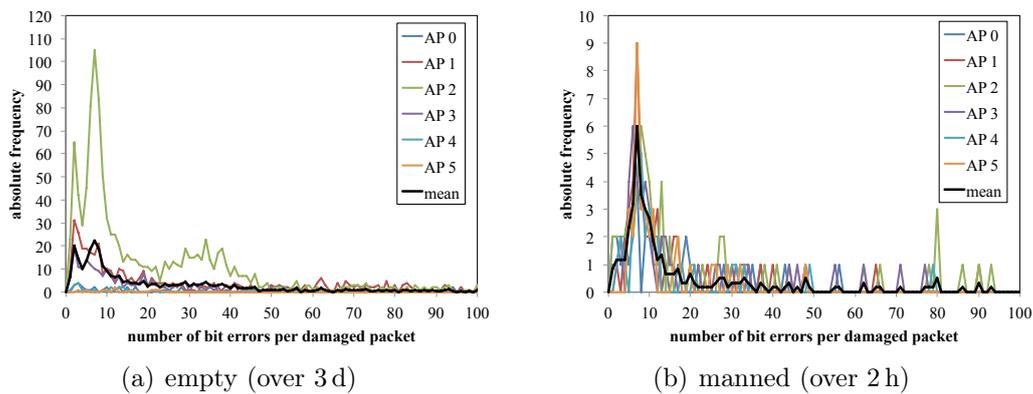


Figure 5.30: Absolute frequency of the bit errors per erroneous packet for both, (a) the empty and (b) the manned scenario.

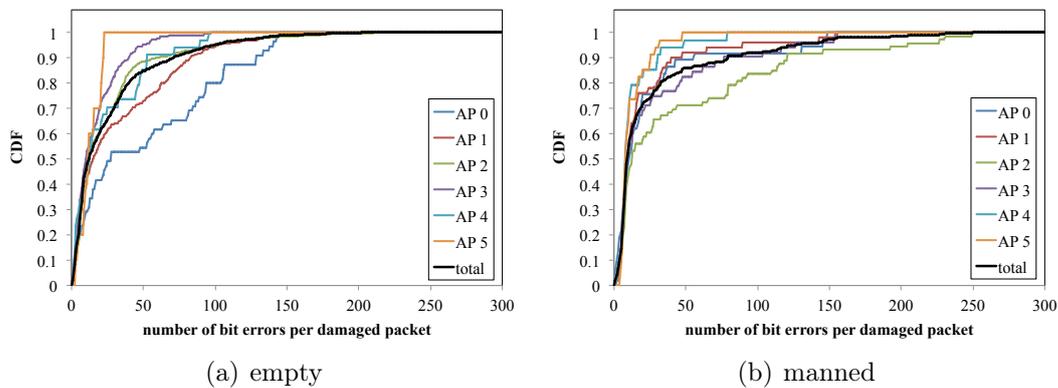


Figure 5.31: CDF of the bit errors per erroneous packet for both, (a) the empty and (b) the manned scenario.

In this measurement campaign only the error detection codes standardized by IEEE802.15.4 were used (a two byte CRC code). But it would be possible to further decrease the system PER when using error correction codes as

well. Error correction codes are able to correct some bit errors but at the cost of increased overhead. Hence, either larger packets are necessary or less available payload per packet is available.

Let us consider the case of block coding forward error correction (FEC) schemes. Several codes are available with different code rates which are capable of correcting a varying number of bit errors. An upper bound for how many bits are reliably correctable in a block code of length l is given by the Gilbert-Varshamov bound. The Gilbert-Varshamov bound states that a block code with k user bits mapped to n channel bits can correct up to t bit errors only if the relation

$$2^{n-k} \geq \sum_{i=0}^t \binom{n}{i}$$

holds. The existence of a code for each triple of (n, k, t) is not assured. However, let us assume that there are codes which fulfill our requirements. Five byte of an IEEE802.15.4 frame are fixed and cannot be included in the FEC scheme. Hence, in our case $n = 528$ bits can be used for the FEC scheme. In order to correct 50% of all erroneous packets the number of correctable bits must be $t = 10$. This would result in $k = 459$ usable data bits. Hence, 13% of the bits would have to be spent for coding. Over 60% of all erroneous packets were received with CRC failures, only 40% couldn't be received at all. If coding each packet with such a FEC scheme which is able to correct up to 10 bit errors the resulting PER could be further decreased by 30%.

Conclusions Mockup Measurements

Spatially distributed access points can eliminate additional errors caused by the presence of humans. The average packet error rate for one access point was around $3 \cdot 10^{-3}$ for the empty cabin but about $4 \cdot 10^{-4}$ for the manned cabin. Using the combination of three access points reduced the packet error rate to its resolution in both scenarios; no packets would have been lost during the measurement time. This clearly shows that with spatial diversity the high volatility of the channels in the manned scenario can be compensated for.

The relatively small number of bit errors in damaged packets means that forward error correction could significantly decrease the packet error rate. However, the exact design of the code must be derived from further evaluations regarding the effect of increased packet sizes due to coding overhead. Additionally, a cross-packet coding scheme could improve the PER even further.

5.6 Conclusions

One of the prerequisites to use WSNs for flight applications is the need for reliable transmissions. Techniques which increase the reliability, as for example Automatic Repeat Request (ARQ), are not feasible because of the already very high network load. Flight applications will depend on very dense sensor networks, hence, all available bandwidth will already be used; retransmissions are no option. Other techniques have to be used to increase transmission reliability.

A custom-tailored TDMA protocol was developed, which introduces spatial diversity by using several access points per wireless cell. The diversity concept exploits the broadcast nature of wireless channels and, hence, can introduce diversity without significantly increasing the consumption of resources. This multi-AP protocol was implemented on specifically design sensor nodes and was evaluated in three measurement campaigns: inside an upper stage of a space launcher Ariane 5, inside an aircraft Airbus A330, and inside an aircraft cabin mockup of an Airbus A330.

It was verified that the resulting packet error rate can be reduced by several orders of magnitude. For most of the nodes, each packet was received by at least one AP during the measurement campaign. And even with passengers inside the aircraft this statement still is valid. This underlines the performance of the protocol to increase transmission reliability for flight applications.

To further increase the transmission reliability a bit error rate analysis was performed to evaluate the performance of supposed error correction codes. It was concluded that by introducing 13% overhead it would be possible to correct up to 10 bit errors and, in doing so, to decrease the PER by an additional 30%.

6 Conclusion & Outlook

In this thesis, it was investigated which challenges may arise when introducing wireless sensor networks for flight applications and what their possible solutions can be. Five major challenges were identified: energy efficiency, transmission reliability, scalability, synchronization, and localization. Out of these five requirements, four were addressed in this work; either by new protocols or by using a wake-up receiver. Table 6.1 gives an overview which chapter addresses which challenges and what means are used to tackle them.

Table 6.1: Addressed challenges for aircraft WSNs and the means to address them

Chapter	WSN demands addressed	Means
Energy-Efficient Clock Synchronization	energy efficiency, reliability, scalability, synchronization	WUR & protocols
Clustering using Wake-up Receivers	energy efficiency, scalability	WUR & protocols
Increasing transmission reliability	energy efficiency, reliability, scalability	protocols

Three protocols were presented which cover various aspects of aircraft WSNs. The first protocol, presented in Chapter 3, covers the challenge of providing a common time base for all communication entities. This has to be done very energy-efficiently for several hundreds of entities at the same time. The presented approach benefits from using wake-up receiver messages for clock synchronization. The power consumption can be reduce by 28 % while decreasing the impact of synchronization packet losses. But to what extent this protocol increases the robustness of the synchronization process against packet losses has to be done in further investigation.

The second protocol, presented in Chapter 4, covers the challenge of reducing the mean power consumption when a lot of sensor readings have to be transmitted. Redundant transmissions will be reduced by clustering sensor nodes on the basis of data similarity. The clustering is done very energy-efficiently by utilizing wake-up receiver messages instead of normal

data packets. Three different methods were presented to monitor cluster validity. Future work should investigate a data adaptive monitoring scheme that automatically chooses the best monitoring method regarding energy consumption and goodput. Additionally, a long-term measurement campaign should be performed to improve the accuracy of the artificial data model.

The last protocol, presented in Chapter 5, covers the challenge of increasing transmission reliability. A TDMA protocol was presented that utilizes spatial diversity using several access points per cell. It was determined that for the particular aircraft scenarios three access points per cell are sufficient to increase the transmission reliability by several magnitudes. This TDMA protocol reduces the PER to such a level that it becomes feasible to use WSNs for flight applications. But some applications still demand even lower PERs, hence, further investigation has to be done. For instance, using maximum-ratio combining instead of selection combining and using a cross-packet FEC coding scheme.

Up to now, the three presented protocols approach different aspects of the challenges that occur when introducing a WSN into flight applications. In the future, these protocols have to be merged to create one common system to tackle all the challenges at once. Additionally, as stated in Section 2.1.5, one prerequisite for introducing a WSN for flight applications is localization, which was not covered by this work. A precise and energy-efficient localization approach has to be included into the system to enable all the possible gains when introducing a WSN.

But when all these challenges are tackled, aircraft manufacturers and airline operators will benefit significantly from introducing a wireless sensor network.

List of Figures

1.1	PERT chart of the thesis structure.	3
2.1	Flight Test Instrumentation (From Airbus S.A.S.)	8
2.2	Structural health monitoring.	9
2.3	Wireless switch for torque limiter indicator.	10
2.4	Ambient Light Control (From Airbus S.A.S.).	11
2.5	Aircraft channel measurements from literature [9–12] together with the free-space path loss model when transmitting with 0 dBm.	14
2.6	An aircraft covered by three cells using different channels. . .	15
2.7	Aircraft channel re-use: (a): path loss off AP 0 and AP 1 separated by 50 m together with the resulting SIR between AP 0 and a receiver, and (b): SIR values between AP 0 and a receiver for several different distances between AP 0 and AP 1.	15
2.8	A 25 % duty-cycled receiver, together with a transmitter that is sending three packets until a successful transmission occurs.	16
2.9	Mean power consumption of duty.cycled AT86RF212 [14] and nRF51822 [15] with different data rate settings.	19
2.10	Power consumption of the low-power receivers using the OOK [23–56] and FSK/PSK [57–67] concepts.	25
2.11	Sensitivity of the low-power receivers vs. their power consumption together with Pareto front.	27
2.12	Gross data rate of the low-power receivers vs. their power consumption together with Pareto front.	28
2.13	From [31]: block diagram of the wake-up receiver from the EU project CHOSeN.	29
2.14	From [68]: schematic of the correlation unit of the wake-up receiver from the EU project CHOSeN.	29
2.15	From [39]: block diagram of the wake-up receiver from the BMBF project AETERNITAS [73].	31
2.16	From [39]: asynchronous oversampling of OOK modulated data bits.	31
2.17	Different communication approaches: (a) when using a duty-cycled transceiver, and (b) when the communication is initiated with a WUR message.	33

2.18	Wake-up delay of low-power receivers [25, 28, 31–33, 37–39, 42–44, 46, 48–51, 54–56, 59, 60, 62, 66] with a power consumption below 200 μ W together with duty-cycled transceivers Atmel AT86RF212 [14] and Nordic Semiconductor nRF51822 [15]; Pareto front included.	34
3.1	Wireless sensor node with its typical components.	38
3.2	From [75]: temperature-dependent frequency characteristic of a tunic fork crystal.	38
3.3	Local times for two nodes; Node 1 with a 10% slower clock ($T'_1 < 1$) and Node 2 with a 20% faster clock ($T'_2 > 1$).	39
3.4	Periodically synchronized local time with a synchronization period of $\delta_s = 30$ s without synchronization errors.	41
3.5	Composition of the delay δ_{total}	42
3.6	Airbus Group Innovations wireless sensor node with WUR.	46
3.7	Jitter of an OV-7604-C7 32 kHz clock [75].	46
3.8	Jitter of the delay between IRQ_{tx} and (a) RF_{tx} or (b) RF_{wur}	48
3.9	Jitter of the delay between IRQ_{tx} and IRQ_{rx} using an AT86RF212 transceiver [14].	48
3.10	Jitter of the delay between IRQ_{tx} and IRQ_{rx} using the wake-up receiver design from Milosiu et al. [39].	49
3.11	Mean power consumption of an AT86RF212 [14] transceiver for varying synchronization periods, together with a theoretical transceiver with very high data rate.	50
3.12	Comparison of the mean power consumption while keeping a maximum clock offset for a duty-cycled transceiver [14] and for a wake-up receiver [39].	52
4.1	Sensor readings from the EU project “ICE: Ideal Cabin Environment Project” [107].	57
4.2	First Markov chain which models the individual drift for each node.	58
4.3	Second Markov chain which models the common group behavior for each node.	58
4.4	Sensor readings of four nodes generated by the data model.	58
4.5	Example for the clustering protocol for three nodes without using a wake-up receiver.	61
4.6	Example for the clustering protocol for three nodes when using wake-up receiver.	62
4.7	Total energy consumption of Node 1 and 31 while establishing the clusters for all 77 nodes with $m = 3$ and $\text{thold} = 2$ when using wake-up receiver.	67

4.8	Total energy consumption of Node 1 and 31 while establishing the clusters for all 77 nodes with $m = 3$ and $thold = 2$ without using wake-up receiver.	68
4.9	Systemwide average of energy consumption per node for the clustering phase over varying numbers of nodes.	69
4.10	Number of cluster leader before an <i>leader announcement</i> was send, the additional leader afterward and the total number of cluster leader; Each with 95 % confidence interval for $m = 3$, $thold = 2$ and $P_{fwr} = 0$	71
4.11	Total number of cluster leader for (a) the missed detection rate and for (b) the false wake-up rate.	71
4.12	Mean power consumption per node and corresponding goodput when using the measurements of the ICE project [107]. Using the conventional as well as the WUR-based approach with all three monitoring methods.	73
4.13	Average power consumption per node and corresponding goodput when using the measurements of the ICE project; including the protocol extension to transmit measurements generated during the clustering process.	74
4.14	Average power consumption per node over reclustering in percent when using the artificial data model.	75
4.15	Goodput over reclustering in percent; with extension to transmit measurements obtained during the clustering process.	76
4.16	Average power consumption per node over reclustering in percent; with extension to transmit measurements obtained during the clustering process.	77
4.17	Comparison of the current consumption of real hardware measurements on a wireless sensor node and the simulated corresponding current consumption of the transceiver.	78
5.1	Multi-AP TDMA superframe structure depicted over time and space.	83
5.2	Mean current consumption of a wireless sensor node over a varying superframe length.	85
5.3	Placement of the six access points inside of the torus.	87
5.4	RSSI values between all combinations of access point pairs.	88
5.5	Signal strength from nodes to access point 0	88
5.6	Signal strength from nodes to access point 1	89
5.7	Signal strength from nodes to access point 2	89
5.8	Signal strength from nodes to access point 3	89
5.9	Signal strength from nodes to access point 4	90
5.10	Signal strength from nodes to access point 5	90
5.11	Highest signal strength from all nodes to one of the six access points.	91

5.12	Highest signal strength from all nodes to one of the three access points 0, 2, or 4.	91
5.13	Setup of access points and wireless sensor nodes in the aircraft Airbus A330.	93
5.14	Message sequence chart for the Airbus A330 measurement campaign.	94
5.15	Average RSSI per node compared to channel measurements from literature [9, 10, 12, 116]	95
5.16	Difference of the mean RSSI values of (a) AP 0 and 1, and (b) uplink and downlink.	96
5.17	RSSI values for AP 0 and 1 in the downlink; (a): RSSI values for Node 165, and (b): RSSI values for Node 45. The diameter of the circles is proportional to the number of measurements.	97
5.18	Category I: stable channel conditions for both APs (Node 485 is depicted).	98
5.19	Category II: stable channel conditions for only one AP; either for (a) AP 0 (Node 410), or (b) AP 1 (Node 16).	98
5.20	Category III: most of the time, both access points have different channel conditions (depicted for Node 332).	99
5.21	PER for all nodes sorted by their PER value for each AP individually for the uplink.	100
5.22	Median of PER for AP 0, 1, and combined.	101
5.23	Airbus A330 Mockup with placement of nodes and APs.	102
5.24	Uplink RSSI between AP 4 and Node 5 for the empty and manned scenario.	104
5.25	PER for all nodes sorted by their PER value for each AP individually. Both, for the (a) empty and (b) manned scenario.	104
5.26	PER for all nodes sorted by their PER for each AP group individually. For both, (a) the empty and (b) the manned scenario.	105
5.27	Mean PER for all AP group combinations. For both, (a) the empty and (b) the manned scenario.	106
5.28	Relative frequency of the burst error length for both, (a) the empty and (b) the manned scenario.	107
5.29	Complementary cumulative distribution function (CCDF) of the run length for both, (a) the empty and (b) the manned scenario.	107
5.30	Absolute frequency of the bit errors per erroneous packet for both, (a) the empty and (b) the manned scenario.	108
5.31	CDF of the bit errors per erroneous packet for both, (a) the empty and (b) the manned scenario.	108

List of Tables

2.1	Wireless application requirements	12
2.2	Characteristics for commercially available transceivers	18
2.3	Requirements for low-power receivers	23
2.4	Some Code Sets of length $(2^n - 1)$ for even n	30
3.1	Jitter characteristic for varying data rates	49
4.1	Simulation Settings	66
5.1	Wireless access point data sheet	84
5.2	Wireless sensor node data sheet	85
5.3	Measurement setup Ariane 5	87
5.4	Measurement setup Airbus A330	93
5.5	Measurement setup	103
6.1	Addressed challenges for aircraft WSNs and the means to address them	111

Nomenclature and Abbreviations

ALC	Ambient Light Control	7
AP	access point	14
ARQ	Automatic Repeat Request	110
ASIC	application-specific integrated circuit	28
ATAG	Air Transport Action Group	1
BER	bit error rate	15
BMBF	Bundesministerium für Bildung und Forschung	30
BPSK	binary phased shift keying	17
CCA	clear channel assessment	83
CFK	carbon fiber reinforced plastic	2
CRC	cyclic redundancy check	82
CSC	conventional selection diversity combining	82
CSMA	carrier sense multiple access	13
ERP	effective radiated power	22
EU	European Union	28
FEC	forward error correction	109
FSK	frequency shift keying	24
FTI	Flight Test Instrumentation	7
GFSK	Gaussian frequency shift keying	17
GPS	Global Positioning System	40
GSC	generalized diversity selection combining	82
ICAO	International Civil Aviation Organisation	13
ICE	Ideal Cabin Environment	56
IF	intermediate frequency	25
ISR	interrupt service routine	45
LiPo	lithium polymer	85

LNA	low-noise amplifier	28
LO	local oscillators	24
LoS	line-of-sight	14
MAC	Medium Access Control	13
MRC	maximal-ratio combining	82
NTP	Network Time Protocol	42
OBS	obstructed line-of-sight	22
OCXO	oven-controlled crystal oscillator	39
OOK	on-off keying	24
O-QPSK	offset quadrature phased shift keying	17
PAQ	Probabilistic Adaptable Query system	55
PER	packet error rate	14
PGA	programmable-gain amplifier	28
PLL	phase-locked loop	30
PoE	Power over Ethernet	84
ppm	parts per million	38
PSK	phased shift keying	24
PWM	pulse width modulation	24
QPSK	quadrature phased shift keying	24
RBS	Reference-Broadcast Synchronization	43
RF	radio frequency	84
RSSI	Received Signal Strength Indicator	86
RX	receive	84
SAW	surface acoustic wave	25
SHM	Structural Health Monitoring	7
SIR	signal-to-interference ratio	15
SNR	signal-to-noise ratio	15
TDMA	time division multiple access	13
TX	transmit	84
UMTS	Universal Mobile Telecommunications System	81
VCXO	voltage-controlled crystal oscillator	40
WBAN	wireless body area network	20
WLAN	wireless local area network	92

WS	Wireless Switch	7
WSN	wireless sensor network	2
WUR	wake-up receiver	2

Bibliography

- [1] ATAG (Air Transport Action Group). Aviation Benefits Beyond Borders. Technical Report April, 2014. URL: <http://aviationbenefits.org>.
- [2] Airbus S. A. S. Airbus Global Market Forecast. Technical report, 2014. URL: <http://www.airbus.com/company/market/forecast/>.
- [3] Johannes Blanckenstein, Jirka Klaue, and Holger Karl. A Survey of Low-Power Transceivers and their Applications. Accepted (IEEE Circuits and Systems Magazine), pages 1–10, 2015.
- [4] Johannes Blanckenstein and Holger Karl. Energy-Efficient Clock Synchronization using Wake-up Receivers. In 22nd International Conference on Software, Telecommunications and Computer Networks (SoftCOM 2014), pages 1–6, Split, September 2014.
- [5] Johannes Blanckenstein, Jirka Klaue, and Holger Karl. Energy Efficient Clustering using a Wake-up Receiver. In 18th European Wireless Conference (EW 2012), pages 1–8, Poznan, April 2012. doi: ISBN978-3-8008-3426-9.
- [6] Johannes Blanckenstein, Jirka Klaue, and Holger Karl. Reducing Redundant Transmissions in WSNs by Clustering using Wake-up Receivers. Submitted (Ad Hoc Networks), pages 1–13, 2015.
- [7] Johannes Blanckenstein, Javier Garcia-Jimenez, and Jirka Klaue. A Scalable Redundant TDMA Protocol for High-density WSNs Inside an Aircraft. In Koen Langendoen, Wen Hu, Federico Ferrari, Marco Zimmerling, and Luca Mottola, editors, Real-World Wireless Sensor Networks, volume 281 of Lecture Notes in Electrical Engineering, pages 165–177. Springer International Publishing, 2014. doi:10.1007/978-3-319-03071-5_18.
- [8] Johannes Blanckenstein, Cristina Nardin, Jirka Klaue, and Holger Karl. Error Characterization of multi-Access Point WSNs in an Aircraft Cabin. In 1st IEEE ICC 2015 Workshop on Dependable Vehicular Communications (DVC), pages 1–6, London, June 2015.

- [9] Genevieve Hankins, Linda Vahala, and John H. Beggs. 802.11ab Propagation Prediction inside a B777. In 2005 IEEE/ACES International Conference on Wireless Communications and Applied Computational Electromagnetics, pages 837–840, Honolulu, April 2005. doi:10.1109/WCACEM.2005.1469714.
- [10] Genevieve Hankins, Linda Vahala, and John H. Beggs. Propagation prediction inside a B767 in the 2.4 GHz and 5 GHz radio bands. In 2005 IEEE Antennas and Propagation Society International Symposium, volume 1A, pages 791–794, Washington, DC, July 2005. IEEE. doi:10.1109/APS.2005.1551442.
- [11] Alexandros Kaouris, Matthaïos Zaras, Maria Revithi, Nektarios Moraitis, and Philip Constantinou. Propagation Measurements inside a B737 Aircraft for In-Cabin Wireless Networks. In IEEE Vehicular Technology Conference (VTC Spring 2008), pages 2932–2936, Singapore, May 2008. IEEE. doi:10.1109/VETECS.2008.306.
- [12] Raffaele D’Errico and Lionel Rudant. UHF radio channel characterization for Wireless Sensor Networks within an aircraft. In 5th European Conference on Antennas and Propagation (EUCAP), pages 115–119, Rome, April 2011.
- [13] Marina Petrova, Janne Riihijarvi, Petri Mahonen, and Saverio Labella. Performance study of IEEE 802.15.4 using measurements and simulations. In IEEE Wireless Communications and Networking Conference (2006 WCNC), volume 1, pages 487–492, Las Vegas, April 2006. IEEE. doi:10.1109/WCNC.2006.1683512.
- [14] Atmel Corporation. AT86RF212 IEEE802.15.4 transceiver data sheet. URL: <http://www.atmel.com/images/doc8168.pdf>.
- [15] Nordic Semiconductor. nRF51822 Bluetooth LE transceiver data sheet, 2013. URL: <http://www.nordicsemi.com/eng/Products/Bluetooth-R-low-energy/nRF51822>.
- [16] Alexandros Elefsiniotis, Dominik Samson, Thomas Becker, and Ulrich Schmid. Investigation of the Performance of Thermoelectric Energy Harvesters Under Real Flight Conditions. Journal of Electronic Materials, 42(7):2301–2305, 2013. doi:10.1007/s11664-012-2411-0.
- [17] Martin Kluge, Dominik Samson, Thomas Becker, Alexander Gavrikov, Benjamin Bennemann, and Joachim Nurnus. Efficient Power Management for Energy Aware, Self-Sufficient Wireless Sensors in Aeronautic Applications. In The 10th International Workshop on Micro and Nanotechnology for Power Generation and Energy Conversion Applications (PowerMEMS 2010), pages 2–3, Leuven, December 2010.

- [18] Benoît Latré, Bart Braem, Ingrid Moerman, Chris Blondia, and Piet Demeester. A survey on wireless body area networks. In Wireless Networks, volume 17, pages 1–18. Springer US, 2010. doi:10.1007/s11276-010-0252-4.
- [19] Stephane van Roy, Francois Quitin, LingFeng Liu, Claude Oestges, Francois Horlin, Jean-Michel Dricot, and Philippe De Doncker. Dynamic Channel Modeling for Multi-Sensor Body Area Networks. IEEE Transactions on Antennas and Propagation, 61(4):2200–2208, 2013. doi:10.1109/TAP.2012.2231917.
- [20] Theofilos Chrysikos, Giannis Georgopoulos, and Stavros Kotsopoulos. Wireless channel characterization for a home indoor propagation topology at 2.4 GHz. In 2011 Wireless Telecommunications Symposium (WTS), pages 1–10, New York, April 2011. IEEE. doi:10.1109/WTS.2011.5960879.
- [21] Hiroaki Tsuchiya. Characterization of the radio propagation channel in residential environment for smart meter communications. In Microwave Conference Proceedings (APMC), 2011 Asia-Pacific, pages 713–716, Melbourne, December 2011.
- [22] Emmeric Tanghe, Wout Joseph, Leen Verloock, Luc Martens, Henk Capoen, Kobe Herwegen, and Wim Vantomme. The industrial indoor channel: large-scale and temporal fading at 900, 2400, and 5200 MHz. IEEE Transactions on Wireless Communications, 7(7):2740–2751, 2008. doi:10.1109/TWC.2008.070143.
- [23] Jose L. Bohorquez, Anantha P. Chandrakasan, and Joel L. Dawson. A 350 μ W CMOS MSK Transmitter and 400 μ W OOK Super-Regenerative Receiver for Medical Implant Communications. IEEE Journal of Solid-State Circuits, 44(4):1248–1259, 2009. doi:10.1109/JSSC.2009.2014728.
- [24] Jia-Yi Chen, Michael P. Flynn, and John P. Hayes. A Fully Integrated Auto-Calibrated Super-Regenerative Receiver in 0.13- μ CMOS. IEEE Journal of Solid-State Circuits, 42(9):1976–1985, 2007. doi:10.1109/JSSC.2007.903092.
- [25] Jeongki Choi, Kanghyuk Lee, Seok-Oh Yun, Sang-Gug Lee, and Jinho Ko. An interference-aware 5.8 GHz wake-up radio for ETCS. In 2012 IEEE International Solid-State Circuits Conference, pages 446–448, San Francisco, December 2012. IEEE. doi:10.1109/ISSCC.2012.6177084.
- [26] Denis C. Daly and Anantha P. Chandrakasan. An Energy-Efficient OOK Transceiver for Wireless Sensor Networks. IEEE Journal of

- Solid-State Circuits, 42(5):1003–1011, 2007. doi:10.1109/JSSC.2007.894323.
- [27] Salvatore Drago, Domine M. W Leenaerts, Fabio Sebastiano, Lucien J. Breems, Kofi A. A. Makinwa, and Bram Nauta. A 2.4 GHz 830 pJ/bit duty-cycled wake-up receiver with -82 dBm sensitivity for crystal-less wireless sensor nodes. In 2010 IEEE International Solid-State Circuits Conference (ISSCC), pages 224–225, San Francisco, February 2010. IEEE. doi:10.1109/ISSCC.2010.5433955.
- [28] Marco Spinola Durante and Stefan Mahlknecht. An Ultra Low Power Wakeup Receiver for Wireless Sensor Nodes. In 2009 Third International Conference on Sensor Technologies and Applications, pages 167–170, Athens, June 2009. IEEE. doi:10.1109/SENSORCOMM.2009.34.
- [29] Patrick Favre, Norbert Joehl, Alexandre Vouilloz, Philippe Deval, Catherine Dehollain, and Michel J. Declercq. A 2-V 600- μ A 1-GHz BiCMOS super-regenerative receiver for ISM applications. IEEE Journal of Solid-State Circuits, 33(12):2186–2196, 1998. doi:10.1109/4.735703.
- [30] Simone Gambini, John Crossley, Elad Alon, and Jan M. Rabaey. A fully integrated, 300 pJ/bit, dual mode wireless transceiver for cm-range interconnects. In 2010 Symposium on VLSI Circuits, pages 31–32, Honolulu, June 2010. IEEE. doi:10.1109/VLSIC.2010.5560262.
- [31] Christian Hambeck, Stefan Mahlknecht, and Thomas Herndl. A 2.4 μ W Wake-up Receiver for wireless sensor nodes with -71 dBm sensitivity. In 2011 IEEE International Symposium of Circuits and Systems (ISCAS), pages 534–537, Rio de Janeiro, May 2011. IEEE. doi:10.1109/ISCAS.2011.5937620.
- [32] Xiongchuan Huang, Simonetta Rampu, Xiaoyan Wang, Guido Dolmans, and Harmke de Groot. A 2.4 GHz/915 MHz 51 μ W wake-up receiver with offset and noise suppression. In 2010 IEEE International Solid-State Circuits Conference (ISSCC), pages 222–223, San Francisco, February 2010. IEEE. doi:10.1109/ISSCC.2010.5433958.
- [33] Xiongchuan Huang, Pieter Harpe, Guido Dolmans, and Harmke de Groot. A 915 MHz ultra-low power wake-up receiver with scalable performance and power consumption. In European Solid-State Circuits Conference (ESSCIRC 2011), pages 543–546, Helsinki, September 2011. IEEE. doi:10.1109/ESSCIRC.2011.6044942.

- [34] Jae Hyuk Jang, David F. Berdy, Jangjoon Lee, Dimitrios Peroulis, and Byunghoo Jung. A wireless sensor node for condition monitoring powered by a vibration energy harvester. In 2011 IEEE Custom Integrated Circuits Conference (CICC), pages 1–4, San Jose, September 2011. IEEE. doi:10.1109/CICC.2011.6055352.
- [35] Pavel Kolinko and Lawrence E. Larson. Passive RF Receiver Design for Wireless Sensor Networks. In 2007 International Microwave Symposium (IEEE/MTT-S), pages 567–570, Honolulu, June 2007. IEEE. doi:10.1109/MWSYM.2007.380553.
- [36] Yao-Hong Liu, Hui-Hsien Liu, and Tsung-Hsien Lin. A super-regenerative ASK receiver with $\Delta\Sigma$ pulse-width digitizer and SAR-based fast frequency calibration for MICS applications. In 2009 Symposium on VLSI Circuits, pages 38–39, Kyoto, June 2009.
- [37] Stevan J. Marinkovic and Emanuel M. Popovici. Nano-Power Wireless Wake-Up Receiver With Serial Peripheral Interface. IEEE Journal on Selected Areas in Communications, 29(8):1641–1647, 2011. doi:10.1109/JSAC.2011.110913.
- [38] Heinrich Milosiu, Frank Oehler, and Markus Eppel. Sub-10 μA data reception with low latency using a 180-nm CMOS wake-up receiver at 868 MHz. In 2011 Semiconductor Conference Dresden, pages 1–4, Dresden, September 2011. IEEE. doi:10.1109/SCD.2011.6068766.
- [39] Heinrich Milosiu, Frank Oehler, Markus Eppel, Dieter Fruhsorger, Stephan Lensing, Gralf Popken, and Thomas Thones. A 3- μW 868-MHz wake-up receiver with -83 dBm sensitivity and scalable data rate. In European Solid-State Circuits Conference (ESSCIRC 2013), pages 387–390, Bucharest, September 2013. IEEE. doi:10.1109/ESSCIRC.2013.6649154.
- [40] Shahaboddin Moazzeni, Glenn E. R. Cowan, and Mohamad Sawan. A 28 μW sub-sampling based wake-up receiver with -70 dBm sensitivity for 915 MHz ISM band applications. In 2012 IEEE International Symposium on Circuits and Systems, pages 2797–2800, Seoul, May 2012. IEEE. doi:10.1109/ISCAS.2012.6271891.
- [41] F. Xavier Moncunill-Geniz, Pere Pala-Schonwalder, Catherine Deholaini, Norbert Joehl, and Michel Declercq. An 11-Mb/s 2.1-mW Synchronous Superregenerative Receiver at 2.4 GHz. IEEE Transactions on Microwave Theory and Techniques, 55(6):1355–1362, 2007. doi:10.1109/TMTT.2007.896796.

- [42] Phillip M. Nadeau, Arun Paidimarri, Patrick P. Mercier, and Anantha P. Chandrakasan. Multi-channel 180 pJ/b 2.4 GHz FBAR-based receiver. In 2012 IEEE Radio Frequency Integrated Circuits Symposium, pages 381–384, Montreal, June 2012. IEEE. doi:10.1109/RFIC.2012.6242304.
- [43] Emil Nilsson and Christer Svensson. Ultra Low Power Wake-Up Radio Using Envelope Detector and Transmission Line Voltage Transformer. IEEE Journal on Emerging and Selected Topics in Circuits and Systems, 3(1):5–12, 2013. doi:10.1109/JETCAS.2013.2242777.
- [44] Seunghyun Oh, Nathan E. Roberts, and David D. Wentzloff. A 116 nW multi-band wake-up receiver with 31-bit correlator and interference rejection. In IEEE 2013 Custom Integrated Circuits Conference, pages 1–4, San Jose, September 2013. IEEE. doi:10.1109/CICC.2013.6658500.
- [45] Brian Otis, Yuen H. Chee, and Jan M. Rabaey. A 400 μ W-RX, 1.6mW-TX superregenerative transceiver for wireless sensor networks. In 2005 IEEE International Solid-State Circuits Conference (ISSCC), pages 396–398, San Francisco, February 2005. IEEE. doi:10.1109/ISSCC.2005.1494036.
- [46] Nathan M. Pletcher, Simone Gambini, and Jan M. Rabaey. A 65 μ W, 1.9 GHz RF to digital baseband wakeup receiver for wireless sensor nodes. In 2007 IEEE Custom Integrated Circuits Conference, pages 539–542, San Jose, September 2007. IEEE. doi:10.1109/CICC.2007.4405789.
- [47] Nathan M. Pletcher, Simone Gambini, and Jan M. Rabaey. A 2 GHz 52 μ W Wake-Up Receiver with -72 dBm Sensitivity Using Uncertain-IF Architecture. In 2008 IEEE International Solid-State Circuits Conference, pages 524–633, San Francisco, February 2008. IEEE. doi:10.1109/ISSCC.2008.4523288.
- [48] Nathan M. Pletcher, Simone Gambini, and Jan M. Rabaey. A 52 μ W Wake-Up Receiver With -72 dBm Sensitivity Using an Uncertain-IF Architecture. IEEE Journal of Solid-State Circuits, 44(1):269–280, 2009. doi:10.1109/JSSC.2008.2007438.
- [49] Nathan E. Roberts and David D. Wentzloff. A 98 nW wake-up radio for wireless body area networks. In 2012 IEEE Radio Frequency Integrated Circuits Symposium, pages 373–376, Montreal, June 2012. IEEE. doi:10.1109/RFIC.2012.6242302.
- [50] Akira Saito, Kentaro Honda, Zheng Yunfei, Shunta Iguchi, Kazunori Watanabe, Takayasu Sakurai, and Makoto Takamiya. An all 0.5 V,

- 1 Mbps, 315 MHz OOK transceiver with 38- μ W carrier-frequency-free intermittent sampling receiver and 52- μ W class-F transmitter in 40-nm CMOS. In 2012 Symposium on VLSI Circuits (VLSIC), pages 38–39, Honolulu, June 2012. IEEE. doi:10.1109/VLSIC.2012.6243778.
- [51] Takahiro Takiguchi, Shunsuke Saruwatari, Takashi Morito, Shigemi Ishida, Masateru Minami, and Hiroyuki Morikawa. A Novel Wireless Wake-Up Mechanism for Energy-Efficient Ubiquitous Networks. In 2009 IEEE International Conference on Communications Workshops, pages 1–5, Dresden, June 2009. IEEE. doi:10.1109/ICCW.2009.5208044.
- [52] Maja Vidojkovic, Xiongchuan Huang, Pieter Harpe, Simonetta Rampu, Cui Zhou, Li Huang, Jef van de Molengraft, Koji Imamura, Ben Bösze, Frank Bouwens, Mario Konijnenburg, Juan Santana, Arjan Breeschoten, Jos Huisken, Kathleen Philips, Guido Dolmans, and Harmke de Groot. A 2.4 GHz ULP OOK Single-Chip Transceiver for Healthcare Applications. IEEE transactions on biomedical circuits and systems, 5(6):523–34, 2011. doi:10.1109/TBCAS.2011.2173340.
- [53] Alexandre Vouilloz, Michel Declercq, and Catherine Dehollaini. A low-power CMOS super-regenerative receiver at 1 GHz. IEEE Journal of Solid-State Circuits, 36(3):440–451, 2001. doi:10.1109/4.910483.
- [54] Toshiki Wada, Masayuki Ikebe, and Eiichi Sano. 60-GHz, 9- μ W wake-up receiver for short-range wireless communications. In European Solid-State Circuits Conference (ESSCIRC 2013), pages 383–386, Bucharest, September 2013. IEEE. doi:10.1109/ESSCIRC.2013.6649153.
- [55] Dae-Young Yoon, Chang-Jin Jeong, Justin Cartwright, Ho-Yong Kang, Seok-Kyun Han, Nae-Soo Kim, Dong-Sam Ha, and Sang-Gug Lee. A New Approach to Low-Power and Low-Latency Wake-Up Receiver System for Wireless Sensor Nodes. IEEE Journal of Solid-State Circuits, 47(10):2405–2419, 2012. doi:10.1109/JSSC.2012.2209778.
- [56] Xiaohua Yu, Jeong-seon Lee, Chang Shu, and Sang-gug Lee. A 53 μ W super-regenerative receiver for 2.4 GHz wake-up application. In 2008 AsiaPacific Microwave Conference, pages 1–4, Macau, December 2008. Ieee. doi:10.1109/APMC.2008.4958320.
- [57] James Ayers, Kartikeya Mayaram, and Terri S. Fiez. A 0.4 nJ/b 900 MHz CMOS BFSK super-regenerative receiver. In 2008 IEEE Custom Integrated Circuits Conference, pages 591–594, San Jose, September 2008. IEEE. doi:10.1109/CICC.2008.4672155.

- [58] Ben W. Cook, Axel Berny, Alyosha Molnar, Steven Lanzisera, and Kristofer S. J. Pister. Low-Power 2.4-GHz Transceiver With Passive RX Front-End and 400-mV Supply. *IEEE Journal of Solid-State Circuits*, 41(12):2757–2766, 2006. doi:10.1109/JSSC.2006.884801.
- [59] Philippe Le-Huy and Sébastien Roy. Low-Power 2.4 GHz Wake-Up Radio for Wireless Sensor Networks. In *2008 IEEE International Conference on Wireless and Mobile Computing, Networking and Communications*, pages 13–18, Avignon, October 2008. IEEE. doi:10.1109/WiMob.2008.54.
- [60] Edward Lee, Phil Hess, John Gord, Howard Stover, and Patrick Necessian. A 400 MHz RF Transceiver for Implantable Biomedical Micro-Stimulators. In *2007 IEEE Custom Integrated Circuits Conference*, pages 173–176, San Jose, September 2007. IEEE. doi:10.1109/CICC.2007.4405706.
- [61] Marten Lont, D Milosevic, Arthur H. M. van Roermund, and Guido Dolmans. Ultra-low power FSK receiver for body area networks with automatic frequency control. In *European Solid-State Circuits Conference (ESSCIRC 2012)*, pages 430–433, Bordeaux, September 2012. IEEE. doi:10.1109/ESSCIRC.2012.6341347.
- [62] Jagdish Pandey, Jianlei Shi, and Brian Otis. A 120 μ W MICS/ISM-band FSK receiver with a 44 μ W low-power mode based on injection-locking and 9x frequency multiplication. In *2011 IEEE International Solid-State Circuits Conference*, pages 460–462, San Francisco, February 2011. IEEE. doi:10.1109/ISSCC.2011.5746397.
- [63] Giuseppe Papotto, Francesco Carrara, Alessandro Finocchiaro, and Giuseppe Palmisano. A 90 nm CMOS 5 Mb/s crystal-less RF transceiver for RF-powered WSN nodes. In *2012 IEEE International Solid-State Circuits Conference*, pages 452–454, San Francisco, February 2012. IEEE. doi:10.1109/ISSCC.2012.6177087.
- [64] Ronald van Langevelde, Michiel van Elzakker, Dave van Goor, Henk Termeer, J. Moss, and A. J. Davie. An ultra-low-power 868/915 MHz RF transceiver for wireless sensor network applications. In *2009 IEEE Radio Frequency Integrated Circuits Symposium*, pages 113–116, Boston, June 2009. IEEE. doi:10.1109/RFIC.2009.5135502.
- [65] Alan C. W. Wong, Ganesh Kathiresan, Chung K. T. Chan, Omar El-jamaly, and Alison J. Burdett. A 1 V wireless transceiver for an ultra low power SoC for biotelemetry applications. In *European Solid-State Circuits Conference (ESSCIRC 2007)*, pages 127–130, Munich, September 2007. IEEE. doi:10.1109/ESSCIRC.2007.4430262.

- [66] Han Yan, Jose Gabriel Macias-Montero, Atef Akhnoukh, Leo C. N. de Vreede, John R. Long, John J. Pekarik, and Joachim N. Burghartz. A 120 μ W fully-integrated BPSK receiver in 90 nm CMOS. In 2010 IEEE Radio Frequency Integrated Circuits Symposium, pages 277–280, Anaheim, May 2010. IEEE. doi:10.1109/RFIC.2010.5477277.
- [67] Rong-Fu Ye, Tzyy-Sheng Horng, and Jian-Ming Wu. Highly sensitive and low power injection-locked FSK receiver for short-range wireless applications. In 2012 IEEE Radio Frequency Integrated Circuits Symposium, pages 377–380, Montreal, June 2012. IEEE. doi:10.1109/RFIC.2012.6242303.
- [68] Christian Hambeck. Ultra-low power wake-up receiver for wireless sensor networks. Ph.D thesis, TU Wien, July 2011. URL: <http://permalink.obvsg.at/AC07811038>.
- [69] European 7th Framework project. CHOSeN. URL: <http://www.chosen.eu>.
- [70] U.S. Goni and Adel M. D. Turkmani. Software estimation of CDMA user capacity. In IEE Colloquium on Spread Spectrum Techniques for Radio Communication Systems, pages 11/1–11/6, London, June 1992.
- [71] Adel M.D. Turkmani and U.S. Goni. Performance evaluation of maximal-length, Gold and Kasami codes as spreading sequences in CDMA systems. In 2nd IEEE International Conference on Universal Personal Communications, volume 2, pages 970–974, Ottawa, October 1993. doi:10.1109/ICUPC.1993.528523.
- [72] Xiangyong Zeng, John Q. Liu, and Lei Hu. Generalized Kasami Sequences: The Large Set. IEEE Transactions on Information Theory, 53(7):2587–2598, July 2007. doi:10.1109/TIT.2007.899528.
- [73] German IKT2020 BMBF Project. AETERNITAS. URL: <http://www.aet-projekt.de>.
- [74] Dawud Gordon, Mathias Berning, Rayan El Masri, Johannes Blanckenstein, Jirka Klaue, and Michael Beigl. WoR-MAC : Combining Wake-on-Radio with Quality-of-Service for Intelligent Environments. In 2012 Ninth International Conference on Networked Sensing Systems (INSS), pages 1–8, Antwerp, June 2012. doi:10.1109/INSS.2012.6240534.
- [75] Micro Crystal Switzerland. OV-7604-C7 clock oscillator data sheet, 2014. URL: http://www.microcrystal.com/images/_PDF/4_Low-power-Oscillator_OV-WM/ov-7604-c7.pdf.

- [76] Maestro Wireless Solutions Ltd. A2100 GPS receiver data sheet. URL: <http://www.maestro-wireless.com/download-dp22>.
- [77] David L. Mills. Internet time synchronization: the network time protocol. *IEEE Transactions on Communications*, 39(10):1482–1493, 1991. doi:10.1109/26.103043.
- [78] Klaus Schossmaier. An interval-based framework for clock rate synchronization. In *Sixteenth annual ACM symposium on Principles of distributed computing (PODC '97)*, pages 169–178, New York, August 1997. doi:10.1145/259380.259437.
- [79] Paulo Veríssimo, Luís Rodrigues, and Antonio Casimiro. Cesium-Spray: a Precise and Accurate Global Time Service for Large-scale Systems. *Real-Time Systems*, 12(3):243–294, 1997. doi:10.1023/A:1007949113722.
- [80] Michael Mock, Reiner Frings, Edgar Nett, and Spiro Trikaliotis. Continuous clock synchronization in wireless real-time applications. In *19th IEEE Symposium on Reliable Distributed Systems (SRDS-2000)*, pages 125–132, Nürnberg, October 2000. doi:10.1109/RELDI.2000.885400.
- [81] Jeremy Elson, Lewis Girod, and Deborah Estrin. Fine-grained network time synchronization using reference broadcasts. *ACM SIGOPS Operating Systems Review*, 36(SI):147, 2002. doi:10.1145/844128.844143.
- [82] Jeremy Elson and Kay Römer. Wireless sensor networks: A New Regime for Time Synchronization. *ACM SIGCOMM Computer Communication Review*, 33(1):149–154, 2003. doi:10.1145/774763.774787.
- [83] Saurabh Ganeriwal, Ram Kumar, and Mani B. Srivastava. Timing-sync protocol for sensor networks. In *First international conference on Embedded networked sensor systems (SenSys '03)*, page 138, New York, November 2003. doi:10.1145/958491.958508.
- [84] Su Ping. Delay Measurement Time Synchronization for Wireless Sensor Networks. *Intel Research Berkeley Lab*, June 2003.
- [85] Marcy Josephine Ammer. *Low Power Synchronization for Wireless Communication*. Ph.D thesis, University of California, Berkeley, 2004.
- [86] Miklós Maróti, Branislav Kusy, Gyula Simon, and Ákos Lédeczi. The flooding time synchronization protocol. In *2nd international conference on Embedded networked sensor systems (SenSys '04)*, page 39, New York, November 2004. doi:10.1145/1031495.1031501.

- [87] Daniela Tulone. A resource-efficient time estimation for wireless sensor networks. In 2004 joint workshop on Foundations of mobile computing (DIALM-POMC '04), page 52, New York, October 2004. doi:10.1145/1022630.1022639.
- [88] Josephine Ammer and Jan M. Rabaey. Low power synchronization for wireless sensor network modems. In 2005 IEEE Wireless Communications and Networking Conference, volume 2, pages 670–675, New Orleans, March 2005. doi:10.1109/WCNC.2005.1424588.
- [89] Alexander Tyrrell, Gunther Auer, and Christian Bettstetter. Synchronization Inspired from Nature for Wireless Meshed Networks. In 2006 International Conference on Wireless Communications, Networking and Mobile Computing, pages 1–4, Wuhan, September 2006. doi:10.1109/WiCOM.2006.244.
- [90] Fikret Sivrikaya and Bülent Yener. Time synchronization in sensor networks: a survey. IEEE Network, 18(4):45–50, 2004. doi:10.1109/MNET.2004.1316761.
- [91] Bharath Sundararaman, Ugo Buy, and Ajay D. Kshemkalyani. Clock synchronization for wireless sensor networks: a survey. Ad Hoc Networks, 3(3):281–323, 2005. doi:10.1016/j.adhoc.2005.01.002.
- [92] Silicon Labs. EFM32LG332 microcontroller data sheet. URL: <http://www.silabs.com/products/mcu/lowpower/Pages/efm32lg-leopard-gecko.aspx>.
- [93] Chunhung R. Lin and Mario Gerla. Adaptive clustering for mobile wireless networks. IEEE Journal on Selected Areas in Communications, 15(7):1265–1275, 1997. doi:10.1109/49.622910.
- [94] Mainak Chatterjee, Sajal K. Das, and Damla Turgut. An on-demand weighted clustering algorithm (WCA) for ad hoc networks. In 2000 IEEE Global Telecommunications Conference (GLOBECOM '00), volume 3, pages 1697–1701, San Francisco, December 2000. doi:10.1109/GLOCOM.2000.891926.
- [95] Ossama Younis and Sonia Fahmy. HEED: a hybrid, energy-efficient, distributed clustering approach for ad hoc sensor networks. IEEE Transactions on Mobile Computing, 3(4):366–379, 2004. doi:10.1109/TMC.2004.41.
- [96] Anand Meka and Ambuj K Singh. Distributed Spatial Clustering in Sensor Networks. In Advances in Database Technology (EDBT 2006), pages 980–1000. Springer Berlin / Heidelberg, 2006. doi:10.1007/11687238_57.

- [97] Daniela Tulone and Samuel Madden. PAQ: Time Series Forecasting for Approximate Query Answering in Sensor Networks. Lecture Notes in Computer Science, 3868:21–37, 2006. doi:10.1007/11669463_5.
- [98] Daniela Tulone and Samuel Madden. An energy-efficient querying framework in sensor networks for detecting node similarities. In 9th ACM international symposium on Modeling analysis and simulation of wireless and mobile systems (MSWiM '06), MSWiM '06, page 191, New York, October 2006. ACM Press. doi:10.1145/1164717.1164768.
- [99] Ameer Abbasi and Mohamed Younis. A survey on clustering algorithms for wireless sensor networks. Computer Communications, 30(14-15):2826–2841, 2007. doi:10.1016/j.comcom.2007.05.024.
- [100] Buğra Gedik, Ling Liu, and Philip S. Yu. ASAP: An Adaptive Sampling Approach to Data Collection in Sensor Networks. IEEE Transactions on Parallel and Distributed Systems, 18(12):1766–1783, 2007. doi:10.1109/TPDS.2007.1110.
- [101] Olutayo Boyinbode, Hanh Le, Audrey Mbogho, Makoto Takizawa, and Ravi Poliah. A Survey on Clustering Algorithms for Wireless Sensor Networks. In 2010 13th International Conference on Network-Based Information Systems, pages 358–364, Takayama, September 2010. IEEE. doi:10.1109/NBiS.2010.59.
- [102] Luca De Nardis, Daniele Domenicali, and Maria-Gabriella Di Benedetto. Traffic-based, mobility-aware clustering for IEEE 802.15.4a UWB networks. In 21st Annual IEEE International Symposium on Personal, Indoor and Mobile Radio Communications, pages 2002–2007, Istanbul, September 2010. IEEE. doi:10.1109/PIMRC.2010.5671592.
- [103] Chao Wang, Huadong Ma, Yuan He, and Shuguang Xiong. Approximate Data Collection for Wireless Sensor Networks. In 2010 IEEE 16th International Conference on Parallel and Distributed Systems, pages 164–171, Shanghai, December 2010. IEEE. doi:10.1109/ICPADS.2010.32.
- [104] Shan-Hung Wu, Chung-Min Chen, and Ming-Syan Chen. Collaborative Wakeup in Clustered Ad Hoc Networks. IEEE Journal on Selected Areas in Communications, 29(8):1585–1594, 2011. doi:10.1109/JSAC.2011.110908.
- [105] Dahlila P. Dahnil, Yaswant P. Singh, and Chin K. Ho. Energy-efficient cluster formation in heterogeneous Wireless Sensor Networks: A comparative study. In 2011 13th International Conference on Advanced

- Communication Technology (ICACT), pages 746–751, Seoul, February 2011.
- [106] Hongbo Jiang, Shudong Jin, and Chonggang Wang. Prediction or Not? An Energy-Efficient Framework for Clustering-Based Data Collection in Wireless Sensor Networks. IEEE Transactions on Parallel and Distributed Systems, 22(6):1064–1071, 2011. doi:10.1109/TPDS.2010.174.
- [107] European 6th Framework project. ICE: Ideal Cabin Environment Project. URL: <http://www.ice-project.eu/>.
- [108] Silicon Labs. EFM32G210 microcontroller data sheet. URL: <http://www.silabs.com/products/mcu/lowpower/Pages/efm32g-gecko.aspx>.
- [109] Real Time Engineers LTD. FreeRTOS. URL: <http://www.freertos.org/>.
- [110] Mathew C. Valenti and Neiyer Correal. Exploiting macrodiversity in dense multihop networks and relay channels. In IEEE Wireless Communications and Networking (WCNC '03), volume 3, pages 1877–1882, New Orleans, March 2003. doi:10.1109/WCNC.2003.1200673.
- [111] Yasuhiro Hirayama, Nobuyuki Nakagawa, Hiraku Okada, Takaya Yamazato, and Masaaki Katayama. Performance analysis of multiple-route packet combining scheme for real-time communications in wireless multihop networks. In IEEE 15th International Symposium on Personal, Indoor and Mobile Radio Communications (PIMRC '04), volume 1, pages 140–144, Barcelona, September 2004. doi:10.1109/PIMRC.2004.1370852.
- [112] Angel Bravo, Antonio Artes-Rodriguez, and Petar M. Djuric. Packet Combining Over Rayleigh Channels using Signal-to-Noise Ratio Information and Detection by the Maximum A-Posteriori Criterion. In IEEE 7th Workshop on Signal Processing Advances in Wireless Communications (SPAWC '06), pages 1–5, Cannes, July 2006. doi:10.1109/SPAWC.2006.346444.
- [113] Muhammad Farukh Munir, Agisilaos Papadogiannis, and Fethi Filali. Cooperative Multi-Hop Wireless Sensor-Actuator Networks: Exploiting Actuator Cooperation and Cross-Layer Optimizations. In IEEE Wireless Communications and Networking Conference (WCNC '08), pages 2881–2886, Las Vegas, March 2008. doi:10.1109/WCNC.2008.504.

-
- [114] Ahmed I. Sulyman and Maan Kousa. Bit error rate performance of a generalized diversity selection combining scheme in Nakagami fading channels. In IEEE Wireless Communications and Networking Conference (WCNC '00), volume 3, pages 1080–1085, Chicago, September 2000. doi:10.1109/WCNC.2000.904778.
- [115] Atmel Corporation. AT86RF233 IEEE802.15.4 transceiver data sheet. URL: <http://www.atmel.com/devices/AT86RF233.aspx>.
- [116] Nektarios Moraitis and Philip Constantinou. Radio Channel Measurements and Characterization inside Aircrafts for In-Cabin Wireless Networks. In 2008 IEEE 68th Vehicular Technology Conference, pages 1–5, Calgary, September 2008. IEEE. doi:10.1109/VETEFC.2008.63.