An Electronic Assistance System for Awareness of Physical Presence

MASTER'S THESIS

submitted in partial fulfillment of the requirements for the degree of

Diplom-Ingenieur

in

Business Informatics

by

Daniel Auferbauer
Registration Number e0625622
daniel.auferbauer@student.tuwien.ac.at

to the Faculty of Informatics
at the Vienna University of Technology

Advisor: Assoc. Prof. Dipl. Ing^{in} Dr^{in} Hilda Tellioğlu

(Signature of Author) (Signature of Advisor)

Vienna
October 19, 2013

Erklärung zur Verfassung der Arbeit

Daniel Auferbauer
Schegargasse 18, 1190 Wien

Hiermit erkläre ich, dass ich diese Arbeit selbständig verfasst habe, dass ich die verwendeten Quellen und Hilfsmittel vollständig angegeben habe und dass ich die Stellen der Arbeit - einschließlich Tabellen, Karten und Abbildungen -, die anderen Werken oder dem Internet im Wortlaut oder dem Sinn nach entnommen sind, auf jeden Fall unter Angabe der Quelle als Entlehnung kenntlich gemacht habe.

(Ort, Datum) (Unterschrift Verfasser)

Abstract (Deutsch)

In Erfüllung der durch die Fakultät für Informatik gestellten Vorgaben für Diplomarbeiten, hier die Zusammenfassung meiner Arbeit in deutscher Sprache: Der Inhalt dieses Textes besteht aus einer Untersuchung der Thematik von „Presence Awareness" als Mittel zur Unterstützung bei der Aufsicht über nicht ortsgebundene Personengruppen. Meine Arbeit stellt eine erste Untersuchung und Evaluation dieses Themas dar, in welchem es meines Wissens nach bisher noch keinerlei andere Forschungsbemühungen gegeben hat. Ziel ist es dabei zu eruieren, ob es überhaupt möglich und von praktischem Nutzen ist, die Aufsicht von Personengruppen welche sich in Bewegung befinden, zu erleichtern indem die Awareness der Aufsichtsperson bezüglich ihrer Schützlinge durch technische Hilfsmittel gesteigert wird. Dabei geht es nicht nur darum die technische Machbarkeit festzustellen sondern auch, unter Einbezug potentieller Zielpersonen für die Nutzung eines solchen Gerätes, die Sinnhaftigkeit des gesamten oben genannten Konzeptes zu evaluieren. Als konkreter Anwendungs- und Testfall wird ein Schulausflug dienen.

Zuerst habe ich eine Anforderungsanalyse des ausgewählten Anwendungsfalles (Schulausflug) erstellt. Zu diesem Zweck führte ich Gespräche mit mehreren Lehrpersonen. Diese Interviews ermöglichten auch einen wichtigen Einblick in die häufigsten Problemstellungen und brisanten Situationen während solcher Ausflüge. Danach erfolgte eine Untersuchung von derzeit auf dem Markt erhältlichen Funktechnologien auf Eignung für das erwartete Anwendungsprofil. Des weiteren beschreibe ich den derzeitigen Stand der Forschung in der Thematik von Reichweitenmessung und Positionsbestimmung in Funknetzwerken. Zur endgültigen Evaluierung des Konzeptes der Unterstützung von Gruppenleitern (durch Assistenz bei der Awareness bezüglich der Anwesenheit aller Gruppenmitglieder) habe ich einen funktionalen Prototypen umgesetzt und sowohl in kontrollierter Umgebung als auch bei zwei Feldversuchen während tatsächlicher Schulausflüge getestet und evaluiert.

Folgendes waren die Ergebnisse meiner Arbeit: Ich klassifizierte drei verschiedene Arten von Situationen, welche für Aufsichtspersonen von Gruppen prekär sind und beschreibe eine Liste von Anforderungen an ein Gerät um diese zu entschärfen. Die Evaluation des Prototypen hat gezeigt, dass das vorgestellte Konzept, die Aufsicht über nicht ortsgebundene Gruppen mit-

tels Unterstützung bei Presence Awareness zu erleichtern, durchaus praktikabel ist. Ich schlussfolgere aus meiner Arbeit, dass ein Gerät, welches diese Idee umsetzt, sowohl die Arbeitsbedingungen von Aufsichtspersonen verbessern als auch die Sicherheit für die Beaufsichtigten erhöhen kann. In der Conclusio wird die Analyse der erzielten Ergebnisse im Zusammenhang mit relevanter Literatur weiter ausgeführt.

Abstract

This document is part of my studies in Business Informatics at the Vienna University of Technology and aims to conclude these by serving as a master's thesis. The subject-matter of my work is a foray into the topic of assisting with presence awareness as a means of aiding in group management. My thesis provides a primary examination and evaluation of this issue – it is, to the best of my knowledge, the first research conducted in this particular area. I aim to clarify if it is possible to help with group management by providing electronic assistance to improve awareness of physical presence and, if so, how this could best be achieved. Not only does this mean an examination of technical feasibility, but also an evaluation of the general concept with prospect users.

The methods and approach employed to that end are thus: I have led interviews with teachers (who are to be considered prospect users) to better understand the problems arising from the management of mobile groups and draw up a requirement analysis. I then investigated wireless technologies currently available on the consumer market and assessed them by a set of criteria, to find the one that is best suited for my needs, as well as delved into the topic of location and range estimation within wireless sensor networks. That being done, I have provided a prototype implementation, which was tested both under laboratory conditions and in the field.

As for the outcomes of my work, I was able to identify three types of problematic situations, which are inherent to the examined use-case scenario, and provide a list of required features as a result of the expert interviews. A prototype featuring basic functionality was implemented and tested. This evaluation has shown the viability of my concept. Most importantly, my results have revealed the technical and practical feasibility of the general subject of assisting in group management by providing technological aid for better awareness of physical presence. I conclude that a device that provides assistance to group management through awareness of physical presence would be of benefit to the working conditions of people tasked with managing mobile groups and the safety of their charges.

Contents

Chapter 1

Introduction

1.1 Problem Statement and Proposed Solution

In our society there exist scenarios that lay on one person the responsibility to ensure the physical presence of some other individuals in its vicinity. This person then has to make sure that all said individuals remain in an area around him or her. The people being monitored in turn sometimes do not care or are not able to stay in that area. Additionally, this whole scenario could be happening while everyone involved is in transit between two locations. Put in a more concise and less meandering way: Sometimes, one person is responsible for a group of people, no member of which should venture too far away.

To illustrate this abstract definition with an example, think of an elementary school class on a field trip. Usually there are one or two teachers acting as supervisors, possibly accompanied by helping hands. They are in charge of between twenty and thirty young children. The teachers cannot allow any of them to venture too far away, lest they lose sight of the others and get lost. This may not seem too daunting a task in a controlled, well accessible environment - yet those are usually few and far between during a school field trip. On its way to the destination of the day's outing, the class often has to navigate busy streets and packed public transports; once they arrive, the group may spread out over an area so wide that a single person cannot keep an eye on all of it.

During a field trip, the teachers are bound to be under a high cognitive load. They have to find their way between locations, navigate the public transport system, catch the right bus or board the correct connection train. They need to make sure that everyone behaves, often talking to the children at length or getting riddled with questions in the process. All the while they still have to assert that not a single one of their charges (of which they have up to thirty) misses the right stop or turns the wrong corner. We all know that children may be easily distracted, that they perceive things differently than their adult supervisors and that they may not always have their mind focused on what their teachers deem important. Two pupils cracking it up at the back of the group can so easily miss a turn or a shout from the teacher. They miss the right tramway stop and off they go; it might be minutes before their supervisors notice they are gone, even when exercising tried routines and best practises. Even when they do everything right, there is a chance that an absentee is not noticed for a few minutes.

To mitigate the teacher's workload during such stressful situations and make them aware of possible stragglers early on, I propose in this work the following solution: a device that keeps track of a group's members and notifies the group supervisor, through auditive or visual signals (or both), if one of his or her charges has gone missing. Said device needs to be mobile which implicates it will be small enough to carry around and runs on battery power supply. The system needs to work both indoors and outdoors and is required to work without any stationary infrastructure, as deployment of such cannot be guaranteed for every location; furthermore it needs to work locally within the vicinity of the supervisor and *not* provide global localisation, so as not to invoke privacy concerns. It shall be easy and quick to use in order to not distract the teacher any further from his or her tasks instead of serving as an aid. The system should be affordable, especially in the light of recent irresponsible budget cuts to the educational system, which will be considered the main area of application in this work. The proposed solution shall scale well at least for groups between ten and forty people, as group size may vary and expansion of the system should be easily achievable to fit needs. I am of the opinion that there is demand for such a device and that it could possibly improve the working conditions of teachers as well as better ensure the children's well-being.

A school class in a setting outside of their class room is my primary incentive for this work. I have often spotted such groups traversing urban environments and wondered at how it is that two people can keep track of so much potential chaos on the move without any support, par maybe some visual aids like making all of the children wear clothing of the same bright colour. With all of the wireless technology currently available on the consumer market, surely it must be possible to do better than that? Hence I resolved to investigate this matter further, in the hopes of manufacturing a prototype which answers that question.

Consequentially to the incentive for this work, a school class on a field trip will serve as primary use-case over the course of this work. However there are also other possible areas of application, which I shall not investigate further in this work. Some scenarios where I could image the proposed device to be of use are:

The led tourist group on its way through a city. I am sure you have seen one of those large groups of people led by a single person in front. Stereotypical this person will hold up an umbrella (though I've really seen them hold the weirdest stuff) so that everyone else can see him or her and follow. People in such jobs could possibly also benefit from the proposed solution, in that it could aid such a tourist group leader by reducing or easing the time he has to "wait up" for his group. It would give him an instant way to see how many people are not with the group any more and what their names are. The "supervisor" probably hasn't seen any of the people he is leading around the city before so it is much easier for him to lose track of his entourage than it would be for a teacher of the primary use-case. This scenario is less challenging in some ways, but more so in others. I would expect it to be easier to implement because the members of the tourist group will probably wear the transceivers more willingly, having understood (hopefully) what they are for. Losing a group member would also be less severe because I expect grown-up human beings to be able to manage on their own, even in a foreign city. Losing sight of the group is an inconvenience, rather than a straight-out hazard. But this scenario also has its own challenges, since the groups are usually larger, with less supervisory personnel. The group leader will rely more heavily on the device to tell him who's there and who's absent.

There is also one possible scenario involving a hiking group. This differs from the two other use-cases discussed until now in that it will take place primarily in the countryside or alpine regions and the group is expected to be smaller. The situation is this: a group of people decide that it would be nice to get out for a bit. They go on a hiking tour through some alpine foothills. The group is about ten people strong, all between twenty-five and forty years of age. The most experienced guy takes the main device and acts as leader, going first. I expect the acceptance of group members towards carrying the transceivers to be quite high since they would have willingly bought and agreed to use the system. The device gains usefulness in this scenario if the group leader is not well acquainted with the hikers. An example for this would be if he was a hired guide. But then this setting also becomes just a variation on the tourist group theme in another environment with a smaller group. The proposed solution could aid the group leader by telling him whether or not everyone is still in range or if he has just lost a member. Even if the group is much smaller than in the previous two scenarios, it can be easy to lose track of a group in rough terrain. With this technical aid, maybe they would not keep going for thirty minutes before they realise an absence.

I also want to lay down which situations the proposed solution is explicitly not intended to be used in. First there are those situations where lives would depend on the correct working of the device. The reason for this is not that I intend to design a system that is unreliable. Rather, it is the case that the proposed solution has different design goals. It is meant as an aid, to mitigate certain inconveniences. They may be strong inconveniences in some cases (like a lost child on a field trip), yet in none of the scenarios described above did I put a person in real danger of losing their lives. It would of course be possible to design a system like that, but factors such as redundancy and backup components for fault mitigation as well as precise localisation would have to receive much more focus than I intend to give them. An example of such a scenario would be a search and rescue operation after an avalanche went down.

Secondly there are circumstances of extremely harsh environments, as a line has to be drawn somewhere as far as the robustness of the proposed solution is concerned. I do not intend to design a device that can withstand the environmental challenges during an Arctic expedition. Any prototype showcased in this work will not be crafted to survive a rafting tour. There will be no support for hikes through the Rub' al Khali desert. More generally speaking the proposed solution is not intended to be used in any activity that may be classified as "extreme sports" or contains words like "expedition" and "exploration".

Lastly there is the case of very small groups. For example, a family of four. There would be two main devices, one for each parent, and two group members, their kids. My concern here is not that the proposed solution would not work. For such a small group the range to cover would be rather small. Creating a prototype to fulfil the requirements set by such a scenario would be much easier than what I have in mind. The problem with this use-case is rather that there exist better solutions for small groups. Examples for such systems would be [18] or [8]. Designing for a scenario like this would allow for different choices and less complex components and software.

1.2 Aims and Structure of this Work

My approach to examine the subject-matter of aiding group management through awareness of physical presence will be thus: First I will lead expert interviews. I have roughly sketched out my proposed solution in the previous section, yet this draft is lacking in detail and a more precise list of required features is in order. Furthermore I do not know if such a system is indeed desired or if it evokes the interest of prospect users. Therefore I will speak directly with members of the target audience, teachers, to get a clearer picture. In this step of the work, I want to get to know the prospect users of my prototype and learn about their wishes and expectations.

Secondly, there will be an assessment of wireless networking technologies that are currently available on the consumer market. Each technology shall be evaluated based on a set of requirements that arise from the proposed solution and the expert interviews. At the end there shall be a justified choice of technology, selecting the one which is best fitted to my needs. Questions to be answered in this part of the work are: What are the currently available wireless technologies? What are the criteria that they should be compared with? Which technology is the best fit for this project overall, with regard to the aforementioned criteria? What are the alternatives and why were they not chosen? This part of the work shall also feature an introduction into the topic of range and position estimation using wireless network technology. I want to find out which, if any, of the techniques that have been discussed in literature are applicable to my project. This section will explain the findings of my research into that topic for the benefit of the reader and later reference. All of the issues named up until now may be found in Section 2.

After having gathered sufficient information on the theoretical parts of this project, the third step shall be to manufacture a prototype. As I am neither experienced in hardware prototyping nor have an in-depth knowledge of wireless networking, no great revelations are to be expected in this part of the work. The prototype shall be a means to an end, and not the main focus. Nevertheless I aim to give an explanation of the construction and programming of all devices built, for those readers that are interested in the topic. The goal of this part of the work will be to provide a proof of concept that it is indeed possible to implement what I have proposed as a solution; because theorising about solving a problem is all well and good, but of little use without a proof that it can be done. Thus there is only one question to be answered in Section 3: is it possible to construct a device that provides an implementation of the problem solution I have proposed? The aim is simply to answer this question with "yes indeed".

The last part of my work shall be an evaluation of the prototype discussed above (assuming I have managed to provide one). First there will come an evaluation under laboratory conditions. This shall serve as an examination of the user interface usability, as perceived by the target audience, and try to answer these questions: Is the device easily understandable and intuitive to use? Are there parts of the user interface that are unclear or confusing? Do the testing participants have questions about the user interface? Is the shape and form of the device pleasing? Severe problems detected during these tests shall be fixed following the closure of the evaluation. The aim of this step is to find and eradicate such parts of the user interface that do not facilitate quick and easy use of the device. After this is done, the device shall be ready for field testing, which is the last part of this project. The prototype will be put to "real" use in a setting that comes as close to its intended field of application as possible. This is meant to answer questions like: Are there any usability issues that did not arise in a laboratory environment? Are there any technical issues that were not detected previously? Does the device affect the behaviour of the participants? Put very concise, this last section is to answer the final and most important question: *does it work?*

The ultimate goal of my work, put in the most concise way I can manage, is this: To examine the possibility and viability of supporting group management through electronic assistance in awareness of physical presence, thereby easing the cognitive load of supervisory personnel and at the same time reducing the risk of stragglers.

1.3 Related Work and Differentiation

To the best of my knowledge, there is very little previous work to be had concerning the topic at hand. There have been many forays into areas that are related to or border on what is described herein, but none that actually tackle the stated problem in a comparable way. I will give an overview over each of these areas that I deem relevant to my work and try to provide reference to some of the work done therein.

There are various commercial solutions available for the purpose of tracking a single child. They provide global positioning services for a single individual, which has to wear a GPS receiver in some shape or form. The information of the child's exact whereabouts is then made available to the parent, most often through a web interface. Examples of such services are [8] or [19]. Studies have been undertaken to investigate the impacts of these or similar on family bonds and trust [32] as well as to why they have not been adopted more broadly [94]. Privacy concerns are noted as major factors in both works. I have also found one commercial product that enables local presence detection [18]. It consists of two devices, one carried by the parent and one by the child, and notifies the former if the latter moves farther than 30 feet away through an auditive signal; in turn it provides the possibility to remotely sound a loud alarm on the child's device. These systems differ from my proposed solution in that they focus on single individuals only. In the case of [8], [19] or simlar, they also provide location tracking on a global scale, whereas I only intend to enable local presence awareness. Their focus is another: providing parents (or other supervisors) with an exact position of their protégée. [18] comes closer to my proposed solution on the technical side in that it provides only local presence awareness, yet it is not a viable solution for groups.

As far as supporting supervising personnel during school outings is concerned, very few attempts have been made, either theoretical or practical. To be more precise, there are only two projects known to me that try to fill this role. One has made the news in 2011 when a daycare centre in Sweden started using GPS to keep track of children. They used transmitters integrated into high visibility vests, the position of which teachers could see on a mobile phone [11] [12]. Here, too, privacy concerns quickly moved to the fore. The difference between this system and the solution I propose is that it could in theory work like the location tracking services mentioned in the previous paragraph. If it is indeed used that way I do not know; yet it has the capability to provide location detection globally and that is something I do not want to provide, as I have said before. Apart from this I am only aware of one other advance on this topic, a patent application called "*Mobile Terminal Group Management System and Mobile Terminal*" [9], which proposes "mutual management" for all group members. Seemingly it is planned to announce departure from the group to both the lost member and all others. Based on what I can garner from the patent application text, this system is meant to provide a much more networked approach than what I have in mind, insofar as all of the devices seem to be concerned with detecting absentees; it empowers all the group members much more through this "mutual management" approach. This is a different concept than my proposed solution in that the latter hands control of the group management solely to the supervisor. Sadly the patent text is not clear or extensive enough to allow any assessment as to whether or not it is to be considered previous work on this topic. Neither have I found an actual implementation of this design nor does the document state any technical details: whether or not it will work locally or globally, which technology it is intended to use or any other specific information.

Despite these ventures into areas related to my work, there have been no publications to date that approach the problem as stated earlier in the same way that I propose in this work. The implementations or patents listed above all differ either in their problem statement or approach to the solution. Of course there is always the possibility that I have missed a publication or have not heard of a commercial product, but it is hardly possible for a single individual to be completely sure to know everything about any topic. Thus I concede to the possibility that I have missed something – yet to the best of my knowledge, there exists no publications to date that tackle the problem as stated in 1.1 in such a way as I intend to.

I feel obliged to mention at this point that the general concept of an electronic assistance system for teachers, which subsequently led me to the more general research conducted in this work, was conceived by a group of students (including myself) during the summer term of 2010; we attended a now discontinued lecture which was titled "Angewandte Innovation". Together we came up with the notion of providing a system that alarms the supervisor if a group member has ventured away beyond a certain distance. The idea for such a device was not mine alone; I feel that fairness dictates I would at least mention this circumstance. Yet apart from the conceptual notion of a local presence awareness aid, this thesis does not contain or reference (let alone claim) any work done back then – the reason being that there is very little substantial material to be had. We merely proposed and presented the idea as part of the lecture. A low fidelity prototype was done based on active RFID tags and a reader for these, which was attached to a PDA. The tags were meant to be integrated into wristbands for the children to wear (which never actually happened). None of the prototyping efforts ever reached a usable state, and work on the project was never continued after the end of the lecture. But since I am picking up the idea now, it has to be considered prior work on the topic.

Chapter 2

Preliminary Inquiries

In this chapter, I will describe my investigations into the theoretical areas related to the project, namely wireless networking, range and distance estimation as well as (to some extent) wireless sensor networks. It provides the foundation for the implementation of the prototype by giving a requirement analysis through expert interviews, providing a preferred technology to use through an assessment and doing an introduction into the topic of estimating ranges and position with radio signals. At the end of this chapter, you will hopefully understand why I chose to implement the prototype as described in Section 3 and how I arrived at that decision.

The chapter is sectioned into three main parts. The first one gives an account of the expert interviews I did with teachers, during which I tried to get to know my prospect users better and learn about their wishes and expectations regarding the prototype as well as their general opinion of the project. I also aim to identify critical situations that occur during school outings, which are of special interest since it is those that I want to mitigate. The second part of this section will be an assessment of wireless technologies that are currently available on the consumer market, based on a set of criteria as defined by my knowledge of any requirements up to this point. At the end I will state a choice and explanation of the preferred technology to use in this project. The last part will be a summary of techniques for range and position estimation using wireless network technology such as has been discussed in literature.

2.1 Requirement Analysis Interviews

As noted in [73, p. 73], "The first step in the usability process is to study the intended users and use of the product". Thus, before starting to actually implement a prototype, I found it would be prudent to talk to some experts on the matter of school outings (and of course, prospect users): teachers.

Interview Goals and Plans

There were several goals to these interviews: First of all I wanted to gauge the reception such a device would get from the supposed target audience. Would there be interest in such a device? Would there be any demand at all for such electronic aids, or would that be regarded as an intrusion into the profession's area of responsibility? It is after all not unheard of for well meant technological assistance to be poorly received by its target audience because "that's not how we do it around here". If the teachers did not want such a device in their daily routines, I would try to find out which features could change that.

Secondly I wanted to get a feeling of the critical areas that typically have to be dealt with during a school field trip. As I had no expertise as far as supervising lots of children is concerned, the aim was to talk to the people who do; about what causes them the most headache when they are out and about with a hoard of young students. Which of the situations that they have to deal with are especially critical, where are the most dangerous areas and where is one most likely to run into trouble?

Lastly I was interested in the expectations of the teachers towards the system I was proposing. If I told them nothing more than the aim of the project, omitting any specifics about the technological background or even the semblance of the devices (because I did not know about those at that point anyway), what would they expect? Which features would they take for granted, what would they wish for in such a system? What would they assume the devices would look like, and what was it capable of? I think that these are strong indicators of the needs of my target audience.

These were the main points I wanted to assess with the preliminary interviews. Some secondary information I wanted to gain were ideas about possible implementations of the child devices. Since the teachers know their pupils well, I assumed that they would have inclinations about which type of apparel would be received well or work best for practical purposes. I also intended to run a list of possible features for the main device by the teachers to see whether or not they liked them or thought them important at the end of the interview.

To acquire some participants for these interviews I utilised a contact inside an elementary school. The teachers there are in charge of children between the ages of six and ten and should provide a good representation of the target audience. I drew up a letter in which I introduced myself and shortly outlined the project I was working on. I furthermore explained that I was working on this as part of my master thesis and how I was looking for people to participate in the project. The purpose of the interview was declared as a requirement analysis and I promised that of course all data would be anonymised. Lastly I offered a copy of the finished work to the participants should they be curious which of their contributions made it into the finished text.

I then gave this letter to my contact at the school. I entrusted her with the selection of participants as she saw fit, only asking of her that she would select from a wide range of age groups. The reason for this was that first and foremost she was familiar with the potential participants already, being better able to gauge who would be apt (read as “talkative”) for an open interview. Second I feared that, if I just hung the letter on a message board in the teacher’s room, there would be too little response; and if there was, perhaps it would be biased to include only those interested in the project to start with, thereby distorting my perception of the whole target group’s interest and demand.

Indeed, a week or so later six persons had declared themselves willing to participate. They ranged in age from being in their early twenties to mid-fifties. There were different backgrounds present both in training and job experience. One participant was currently in charge of a pre-school class of very young children, two did also have experience as kindergarteners. Due to the difference in age, their form of qualifications also differed (along with their experience on the job, of course). Overall the pool of participants felt quite divers, and, by being of the target audience, the best representation of a potential user I could get my hands on.

The interviews would be held at the workplace of the participants. Since I wanted to get to know the user, it seemed only fitting to do so at the location of their daily jobs. I would make individual appointments with each participant, which were chosen such that the children of the class had already left, but not so late that the teacher would have to remain at school just for the interview.

I decided to lead a semi-structured interview. There were certain topics towards which I wanted to steer the interview, but I did not intend to squeeze it into some tightly structured questionnaire. I wanted to get to know the user, after all, and this means having them talk and tell as freely as possible about their work. To prepare for the interview, I wrote down a list of topics which I wanted to discuss at the very least, everything else I left up to the flow of the talks. These topics were to be included in every interview:

Problematic Situations What are the most difficult and stressful situations typically encountered on a school outing? Where is it especially difficult to retain an overall view of the children, or put differently: where are children prone to get separated from the group and lost?

Best practices What are the individual best practices and habits that the teachers employ to keep the children together and make sure nobody gets lost? How do they deal with situations like they were discussed in point one? What works best, what has proven less efficient?

Known cases Descriptions of cases where children have actually been lost on a school outing, may be witnessed first hand or heard from a colleague. How do you think this could have been prevented? What, if any, are the emergency plans if a child gets lost?

Desirability If I offered them a device that would help keep track of the children, would they welcome such aid or think of it as just another electronic burden to carry around? What would they expect from such a device? Which features would they want, if they could write a wish list with total disregard of technological feasibility?

Feature proposal I describe a list of features that I have come up with (if they have not been mentioned in the talks up until now) and ask the participant how useful or viable he finds them.

Before the start of each interview I would take time to talk to the participants about the project and give them a short introduction. I would explain what the goal of my work was, why I was doing these interviews, what place they had in the final work and how long they would approximately take. I explained that I did not have an extensive questionnaire with me, but rather just wanted to chat with them about their work and school field trips. I assured the participant that he or she would not be called by name in the final work and that everything would be kept anonymous. Lastly I asked them for permission to take an audio recording of the interview session and asked whether they had any lingering questions concerning the interview or the project.

After that I would start the audio recording. The interviews took an average of fifteen minutes each. The talks proved to be quite relaxed with most of the participants talking freely and seeming at ease. The interviews provided some interesting insights into all of the my main areas of interest.

Three Problematic Situations

On the topic of the most problematic and stressful situations during a school outing, there were three scenarios to be identified, which were each addressed by multiple participants. After all the interviews were done I consider these to be the areas were teachers would greatly benefit from electronic assistance in keeping track of the children.

First of all there was the entering and exiting of public transport vehicles. This topic was mentioned as being very difficult and stressful by all of the participants. The situation is as follows: the group waits at some public transport station for a train or a bus. As soon as the vehicle arrives the teachers have to make sure that all of the children enter and none are left behind on departure. Similarly, as soon as they arrive at their destination, everyone has to exit again as fast as possible, lest one of the children involuntarily continues the ride on its own. Since there are a lot of people involved in both cases, it gets very difficult for the teachers to keep an eye on all of their charges. Most participants stated that exiting vehicles is more demanding than boarding. This is due to the fact that transports are often crowded, and school children, being small by nature, tend to get overlooked.

The best practices used to deal with these situations seem to be common throughout the school: teachers usually work in pairs when doing school outings, so when at a station it is customary that one takes the lead and the other remains behind, to make sure that none of the children are left outside before entering, too. The same thing is true in reverse when exiting a transport. After the group has gotten off the vehicle, the teacher tries to bring some order into things, the children line up in pairs of two. Only after this is accomplished can the teacher start a counting of heads. By this point, the transport has usually left the station already, so any missing children remaining aboard would most likely be discovered too late.

Three of the six participants also remarked that they usually have some form of emergency plan in place with the children, should one of them get lost. This usually amounts to telling their charges that, if they should miss the right stop and remain on vehicle, they should travel on to the next station and wait there for someone to pick them up. In addition to that, one teacher said she has taken to handing out cards with her phone number to the children so that they could ask someone to contact her in case they are lost in transit.

Some explanations as to why a child might get lost in public transport at all were that, for one, it is often difficult to communicate with children inside vehicles. Those are often crowded with passengers, and it could well happen that the children do not perceive the teacher's call to get off at the next station. Then there is the factor that, in crowded places, people will generally behave erratic and even less considerate than they are otherwise; this behaviour does not make exemptions for children, so sometimes they are just being blocked from exiting. Lastly there is the time constraint: even if all goes well there is the possibility that the operator of the transport will close the doors before all of the children have boarded or exited.

I think that I could actually provide some relief as far as the first point, not hearing the teacher's signal to get off now, is concerned. It should not be too difficult to outfit the children's devices with some form of notification system, either through sound or maybe visuals, like a coloured LED light at a highly visible position. The main device would then have an option to notify all the children in the group. This would surely be useful as an alternative to shouting all across the transport, something which two of the participants remarked to be uncomfortable with anyway.

The second problematic situation was identified to be a stay on wide but badly assessable grounds. This scenario type was mentioned and emphasised only slightly less than the public transport one. This point concerns situations were the whole group does a prolonged stay in an area were the children are allowed to roam freely. Examples for this are playgrounds (which are part of many school outings during warm seasons), fairs and markets (especially popular around Christmas) or gathering leaves and plants in the forest (for school projects). In these situations, the children are supposed to be allowed freedom of movement to go about their business so long as they do not leave the designated area.

The most important factor here, what makes these situations so precarious, is that the teachers are prone to lose visual contact with their charges. They may hide behind trees, turn a corner while playing, get lost in a group of people or similar; there's a myriad of ways for a child to disappear out of sight. There is little in the way of experience or best practices to remedy this fact. Some of the teachers stated that though they would want to, they do not let their class roam free at such occasions, instead trying to maintain structure within the group. Others who let the children disperse throughout the area said that they do not really feel comfortable with it, but there is nothing much to do apart from trying to imbue upon the children the importance of not venturing too far and then trust them.

During the stay the teachers sometimes remain in one central position, which serves as a point for meeting and gathering in the area. Children come by there to ask questions, eat lunch or rest. If the children are younger or the class is less reliable, the teachers tend to wander the area to make sure there is no trouble. If there is enough supervisory personnel available (e.g. if some parents accompany the class to help out), the class may also split up in small groups during their stay, thus granting the children some freedom of movement while some form of oversight is still maintained.

Those participants who did not let their group disperse in these open areas stated that this was out of necessity, not because it is their preferred way. They did not dare to let the group's structure slip for fear of losing a child. I am convinced that an assisting device such as the one that I am proposing could elevate their fears and allow them to grant their charges some additional freedom.

However, most of them also seemed to expect (or wish for) some form of locating technology to be implemented, so that they could find a child quickly in case it gets lost. Alternatively there was some desire for the possibility of making direct contact with the pupils through the device or a notification on the child's device that would warn them if they ventured off too far. I have to admit that recognising the absence of a child only meets half of the requirements in such a situation. To really be of use, there would have to be some way to locate the child, or communicate with it to give instructions for a safe return. Otherwise the teacher will have to search all of the area to find his or her missing pupil. Notifications about children being out of the allowed range would definitely be a step in the right directions, but it would not diffuse the situation completely.

The last problematic category I could identify concerns in-transit passage of highly frequented areas. This includes, among others, crossing streets at traffic lights or changing public transport at a junction. The problems arising in these situations are in some aspects similar to those of the first scenario discussed in this section: they involve lots of people that behave as we have come to expect in large groups where everyone has some place else to be (read as: inconsiderate with a touch of tunnel vision) and a certain time constraint. Added to that is the factor of a moving group.

This problem was mentioned by two of the six participants. The difficulty arises from the fact that the structure of the group is often disrupted by people. People sometimes walk trough the group without regard for the children, who may get confused and loose their group in the crowd. This is especially dangerous at junctions (e.g. of subways) where a lot of passengers are heading in different directions within the same area. The teachers voiced concerns that their charges may get blocked off from the rest of the group, lose their bearing and get lost.

Another example is the crossing of a street at a traffic light. The group has to get across as fast as possible while they have a go from the traffic light. Together with them there is usually a bunch of other people crossing at the same time. These may also disrupt the order of the group, such that a child might be left behind at the wrong side of the street or take a wrong path afterwards. It seems easy to loose such a small individual amidst a rush of pedestrians.

Situations that seem to exert less strain on the supervisory personnel, on the other hand, are those in enclosed areas where a structured procedure is being followed. Specifically, one participant mentioned museums as an example for somewhat relaxed situations. Another example would be a guided tour through some site. In such cases it is less likely to lose visual contact with a child, to miscount or to run into time constraints.

Wishes and Expectations

The teacher's expectations concerning the device's features that I could gleam through the interviews varied somewhat. They did however feature a common core, a set of functionalities, each of which was mentioned by at least four out of six participants. Without me giving them much information about the device that *I* had in mind (because I wanted to find out what *they* did, instead), the participants seemed to have fairly low expectations as far as the features of the device were concerned, mixed with some that are very hard to implement, at least for somebody with my knowledge of radio transmissions and wireless networking.

The less challenging requirements were mostly congruent with what I had had in mind as prototype features anyway; they form the bare minimum of what the prototype should be capable of. The core functionality should be that the device sets off an alarm if a child is separated from the main device, no surprise there. Furthermore the system should display the name of the pupil in question. These were the core features that all of the participants agreed upon.

There were also several mentions of how some form of group member management would help. When questioned about this the participants stated how it is sometimes difficult to remember which of their pupils was present at the start of the outing, who was ill and who has actually gone missing. To better keep track of this, they requested the possibility to enter the present group members by name in the morning so that the list would be saved for the rest of the day.

Among the wishes expressed by a majority of the participants there were two that I judged to be very hard to implement: the first one was *immediately* sounding an alarm if one child was separated from the group. A child remaining inside the tramway while the rest of the group has exited or taking a wrong turn at a junction were two example scenarios mentioned for when the device should immediately display a warning. Therefore this requirement is imminent to the first and third problem scenarios, public transport and crossing highly frequented areas. This is a difficult thing to implement for me because it requires very fine grained resolution of a node's distance to the main device or at least its neighbours; I would however try, to best of my abilities, to make it happen for the prototype. Please refer to page 70 and following for more information on that (ultimately failed) effort.

The second hard to implement feature that was requested by multiple participants was some form of location service for the main device. As was stated in the previous section, realising that a child is missing is often just the start. In many environments, it would be difficult for a teacher to track a child that he has lost sight of or contact with. Naturally, the possibility to pinpoint the position of the pupil in question would be an invaluable help. Sadly this is something that I knew I would not be able to do since the theoretical spadework was done. I explain why on page 40 and following.

A new notion which I had not previously entertained was that of communication from teacher to pupil via the system. This was wished for by several participants, as it would make their job of indicating to the children that an action on their part was required much more easy. As I have stated on page 12, there are situations in which a teacher may have difficulties communicating with the whole group. Examples of this are signalising to exit the vehicle in a tightly occupied public transport or prompt a gathering at one central point, while the children are dispersed throughout a wide, poorly visible area. For such cases, many participants wished for a feature to give a signal to their group.

I was of the opinion that this feature could very well be implemented in my prototype. Though I did not yet know what exactly the child devices would look like in the end, I resolved to include some buzzer and LED lights in my next order of electronic parts. I judged that one way communication with some visual or auditive signal on the child's end should not be too much additional work.

Though I tried to convey that the participants should not be concerned about what they think was technically realisable when stating their expectations for the device, most of the chose to remain rather conservative in their remarks. There were very few requests that ventured beyond the features discussed in the above paragraphs. There was one mention of a two way voice communication between teacher and pupils, like one would have with a handheld receiver (a "walkie talkie").

The aim of this feature was twofold: firstly to determine a child's position through talking, in lieu of a positioning system. Secondly to be able to calm an upset child that has lost contact with the group. Undoubtedly this would be of great use for the teacher. Nevertheless, implementing such a thing will be beyond the scope of my prototype, as it would require both microphones and speakers in all of the devices.

Another feature proposed by a participant was a "panic button" on the child devices. This was meant to give the child some way to notify the teacher if something went wrong. When questioned about whether she expected the children to abuse such a system, she remarked that they most likely would not do it out of malevolence, but rather curiosity and playfulness; maybe if they were allowed to play around with the feature at the very beginning to test it out, they would not activate it without need afterwards. Nevertheless, the teacher did not seem sure that the whole class could be relied upon to have the discipline necessary for being handed such a system.

Various and Assorted

Some participants gave insights into what they thought would work well as designs for the child devices. One remarked that vests would be the best form of apparel to use, since they were versatile enough to be worn in both warm and cold weather, above a child's other attire. Specifically, she proposed high visibility vests for this purpose. I liked this idea very much and would actually later pick it up and implement such textiles for the field tests.

Another teacher remarked how, in her experience, boys preferred wristbands and girl necklaces when it comes to accessories after I mentioned how I had thought about implementing the transceivers in such apparel. This, too, was an idea I would later implement by making one of the textile prototypes into a pendant mainly targeted at girls. Wristbands were also recommended by a second participant, who furthermore reacted doubtful to my idea of implementing the child devices as caps. In her opinion, headgear is too easily taken off and thereby prone to being discarded by the children. I noted this down as interesting and to be examined during field tests.

Concerning the form of the alarm signal, most participants would prefer audio signals in a short, succinct manner which could also be heard in loud environments. They need to be distinct enough to be immediately recognised as the device's alarm and not mistaken for another device's notification (such as a mobile phone ringtone). I also asked about the viability of a silent option that includes vibration, like it is common with mobile phones; most participants stated that this was of small use, since they usually did not take notice of such forms of notification and felt they were too easily missed, especially if the device was carried in a bag.

I also asked the teachers whether or not they thought that specific range data would be of use. Fore example, if the device indicated that a child was 70 metres away or if they could choose the allowed range in ten metre steps, would they be able to relate to such information in the heat of wresting control over a group of children? As it turns out, the majority of the participants did not think that they would gain any benefit from exact distance information as opposed to more fuzzy range definitions. The reason for this was always that the teachers were not sure they could accurately judge the distance of a given number of metres. This is made even more difficult in stressful situations. Therefore I resolved to implement fuzzy range categories along the lines of "near" and "far".

Conclusion

To conclude this section, I will give you a short summary of the results provided by these preliminary interviews. The reception of the general idea of a device that assists in group management by providing information about absentees was received very well. The participants all seemed genuinely happy that somebody recognises the difficulties they have to deal with on school outings. All of them gave the impression of being interested in the project; four out of six asked to try out the prototype on their own initiative. I infer from the overall reactions that acceptance of the device would not be a problem at all. On the contrary: If my selection of participants is as representative of the target audience as I think they are, then there is quite a demand for such a system.

When defining the interview goals I also resolved to find areas and situations in the course of school outings that are especially taxing on the teachers. Three problematic scenarios were identified where a supervisor is prone to lose track of his or her charges and a child might be separated from the group. The common problem is that, to take stock of possible absentees, the teacher has to make the whole group stand in rank and file to do a count of heads. This takes time and is also impossible (or rather, impractical) to conduct in a continuous matter. Therefore, it could be minutes before an absence is discovered. Providing a system that does the job of checking present group members in very small intervals will surely remedy this, at least to a certain degree.

The last of the main goals was to discover what features the participants would expect of a device such as I proposed. I found that the basic requirements seemed to be fairly low: they merely included a list of pupils and a warning if one of those ventured too far from the group, including a message that indicated which one was the child in question. Some wished for a feature to better communicate with the group, at least one-directional (from teacher to pupil). This is an interesting point I had not previously thought about and would surely try to implement. Then of course there were the more problematic wishes for localisation of all group members. Admittedly, this would solve a lot of problems, but is also very difficult to implement.

Here is a short listing of the feature categorisation as found out through the preliminary interviews. *Core* includes those features that are absolutely necessary for a valid test of the device, *extra* lists functionalities that were indicated by the participants to be nice to have, but which I judged to not be indispensable. Lastly there is the *Ambitious* category, which consists of feature requests that were noted but may prove too difficult to implement over the course of this work.

Core Detection of complete absence from the group (meaning the device is no longer part of the network); in case of an alarm, display the name(s) of the missing group member(s); give a concise audio signal that is immediately recognisable.

Extra Provide some form of group management functionality: set up a list of members at the start of the day, which is afterwards used as basis during the trip; being able to choose a maximum distance which the children are allowed to venture from the main device on the fly; sound an alarm not only if a child is completely lost, but also if it diverged more than the chosen maximum distance from the main device.

Ambitious Locating child devices, either their absolute positions on a map or relative to the teacher's current whereabouts.

Apart from the above I also had some interesting conversations about what the teachers think would work well as design for the child devices. This yielded some interesting ideas, like high visibility vests, to try out during field testing. Furthermore I discarded my idea of implementing a silent alarm based on vibrations, as none of the participants seemed to be particularly taken with the idea. After the requirement analysis interview I did indeed have a clearer picture of what exactly the problems are that were to be solved.

2.2 Assessment of Technologies

There exists a variety of technologies on today's market that have the capability to serve as a basis for the prototype envisioned in Section 1.1. In this chapter I will give an analysis of the different technologies that are available. First I will present and explain several criteria by which to assess the technology. Following that I will introduce a number of technologies and measure them up against my criteria. At the end of the section I will present my conclusion and justification of choice.

Assessment Criteria

Overall I have defined five criteria by which to measure wireless technologies. These are:

Range To be of practical use my device will need to cover a certain range. By that I mean it will have to maintain an area of influence within which it is able to detect the presence of members of its network. To achieve this it will be necessary for the underlying communication technology to provide a certain minimum range. It is difficult and maybe not all that beneficial to put in an explicit number here. According to the preliminary interviews, the range within which children may move around the teacher extends from thirty metres during transit to more than a hundred metres during a stay on e.g. some playground. For now, I will assume a minimal range of 50 metres to be required. I dare say even before I have done any assessment of the available technologies that it will be difficult enough to achieve a stable connection over more than 50 metres in an urban environment without a truckload of amplification equipment.

Interference Resilience I expect there to be an abundance of possible interferences to the operation of my system. These will be twofold: first there is the possibility of interferences from other radio sources. Of the five technologies I am going to examine in this chapter for example three operate in the ISM frequency bands. That is a harsh environment in any urban area due to the sheer amount of systems operating there. So there is a need to mitigate interferences from other radio transmissions.

> **ISM bands - the Crowded Space**
>
> The ISM bands are a set of unlicensed frequencies that are reserved for **I**ndustrial, **S**cientific and **M**edical use. Among others there are 868 MHz for Europe, 915 MHz for the Americas, and 2.4 GHz for worldwide use [93]. Of these the 2.4 GHz frequency band (ranging to approximately 2.5GHz) is the most relevant to this work due to its worldwide availability. Popular and widespread occupants of these fields are WiFi (page 25), Bluetooth (page 22) and ZigBee (page 33). Apart from these well-known and, more importantly, standardised protocols there is a numberless horde of devices utilising these frequencies. An example for this that probably does not come to mind immediately are microwave ovens. They operate at or around 2.45 GHz. In theory they should shield their environment from leakage, nevertheless it still occurs [50]. There have been extensive surveys as to how microwave ovens affect wireless communication within the ISM bands (see e.g. [54] for WiFi or [50] for ZigBee). This is just an example of course, but it illustrates just how wide the range of devices utilising the green pastures of ISM bands is.

Secondly there is interference through physical obstacles. Think of walls in a building, trees and boulders in an outdoor environment or practically everything on a busy city street. All these things pose interferences to a radio signal both through the signal strength falloff that occurs when they are penetrated by the transmission and by the multipath problems they induce on the receiver's end. So the ability of a technology to penetrate through these will also matter. This pretty much boils down to what frequency it operates on.

To sum it up short I am going to examine how reliable a technology is able to get a transmission from sender to receiver in the face of interferences. The reasoning behind this is that a system that gives false alarms every few minutes will hardly induce any kind of trust in the user.

Availability The dissemination of the technology on the consumer market and its general availability should not be underestimated as a factor. It may very well be that there is a technology out there which fulfils all the needs of this project to the letter. This will however be of little use to me if said system is some military grade device currently locked away in a bunker. On one hand a widely used and accepted system will make it easier to facilitate a prototype. There will have been prior research on the technology and the necessary components will be more readily available to me. On the other hand good availability and dissemination on the consumer market means that a technology will run a much smaller risk of becoming disused or obsolete soon. This in turn reduces the probability that I will have wasted time, energy and money on doing research based on a soon-to-be dead technology.

Infrastructure Dependencies The technology I am looking for should be free of dependencies to external infrastructure or third party services. By this I mean that it should be possible to establish a connection between two devices without the need for a fixed infrastructure. An example of such infrastructure would be the radio towers of the GSM network. The reason to avoid fixed infrastructure is simply that I cannot provide it. There has been an ample amount of research into sensor networks with a fixed infrastructure, yet the groups and individuals sketched out in my use-case scenarios are all inherently mobile. Counting on a fixed point infrastructure will severely limit this mobility, making my device essentially useless. Therefore it has to work without relying on the comfort of having stationary reference points. Put briefly: I want my system to be as independent as possible.

Energy Consumption Long battery lifetime is not an optional feature. The system needs to remain operational for as long as possible without having to recharge. This is to guarantee that it remains operational through each of the scenarios listed in Section 1.1. Several hours of operation is something that I consider a lower bound here. Ideally the system should remain operational a few days without the need to renew energy sources. Thus I am looking for a technology that enables energy preservation to a point which allows for the prototype to run at least six hours without recharging.

Having introduced the criteria by which to measure, I am now going to discuss each of the technologies and my assessment of them in turn. The order in which they are listed is arbitrary and does not reflect preference. My conclusion and choice of best fit are at the end of this section (page 38).

First Technology: Bluetooth

Bluetooth is a wireless technology for close range personal networks operating on the ISM frequency band and aimed at replacing physical cables.

The name first came up in 1994 when Ericsson Mobile Communications investigated alternatives to the well established concept of physical cables. As Jennifer Bray tells us in [35], the name is taken from Harald Blatand, a tenth-century Viking known to historians (and seemingly some people in the electronics and informations technology industry) for uniting and ruling Denmark and Norway. This was supposed to be somewhat of a good omen for the Bluetooth technology to unite devices via wireless communication. Indeed it has become a major player in wireless communications, though not the one technology to rule them all. Bluetooth's standard and development is overseen by the Bluetooth Special Interest Group. This body comprises of corporations that use, develop and push the technology; notable members include among many others Intel, Microsoft, Apple and Motorola. To use the Bluetooth trademark a company must first become a member of this Special Interest Group.

Range The specifications of Bluetooth name three different classes of radio power. These have not changed since the technology's inception. They are depicted in table 2.1, as described in [35]. Note that the distances are *maximum* distances. It is likely that they were tested in favourable laboratory environments and amount to less in real world scenarios (especially urban ones). Even if taken with a grain of salt 100 metres is a respectable range as wireless communications go. A downside to this is inherent in the design of the Bluetooth technology: Bluetooth is intended as a replacement for physical wires, supports video and audio and is bidirectional. This means that a 100 metre link requires a class one device at *both* ends.

Class	mW	dBm	range (m)
1	100	20	100
2	2.5	4	25
3	1	0	10

Table 2.1: Bluetooth's radio power classes

Apart from the obvious maximum distance Bluetooth also has a minimum distance, stated in [35] to be about ten centimetres. A smaller distance could make the communication unreliable. While this is noteworthy it is unlikely to negatively influence the functionality of my device. For a group member's beacon to be within ten centimetres of the group leader they would have to be standing on each others toes; someone standing on your toes is a pretty surefire way to know that you have not lost them.

Interference Resilience Bluetooth uses the ISM bands (located at 2.4 GHz) which are also occupied by numerous other technologies - among others WiFi (2.2) and ZigBee (2.2). The technology's way to cope with this crowded environment is frequency hopping: to avoid degradation or outages in transmissions two Bluetooth device's communication with each other will use a unique, pseudo random sequence of channel switching. 79 of such channels are available, each with a width of one MHz. The sequence is derived from the master device's address and a clock value; each slave device in the same network synchronises their channel hopping to this sequence. The interval is 1600 hops per second. In [39] it is stated that for reliable data transfer rates it is necessary to keep WiFi (802.11b, specifically) and Bluetooth at least three metres apart. In [91] Shuaib *et al.* show that Bluetooth performance is influenced (sometimes heavily) by the presence of an IEEE 802.11g device. The throughput of the Bluetooth device suffered a reduction of between 17% and 36% in their test cases. According to [78] Bluetooth possesses non-collaborative methods (i.e. not depending on the other device's cooperation) to mitigate blocking by WiFi. They all use similar approaches to detect traffic and declare certain channels as "bad" if they suffer from interference. It is also argued in [78] that these methods are not always effective in practical situations, sometimes starving the Bluetooth device due to excessive courtesy of the transmission protocol. But despite the hostile environment that Bluetooth operates in it seems to be reliable enough for the needs of this project. Since the user does not necessarily need to feel a high packet loss rate (unlike say in audio transmission) and indeed a very low data transmission rate will be necessary, Bluetooth communication can be expected to be stable enough.

Dissemination and Availability As I have said in the introduction, Bluetooth had quite the dashing start. In 2001 Bray called it „*The fastest growing technology since the Internet*" with an estimation of 200 Million BT-enabled devices shipped that same year - only three years after publication of the standard [35]. According to the Bluetooth Special Interest Group's annual report seven billion Bluetooth devices were in market by the end of 2011. 1.8 billion devices were being shipped that year alone [5]. Of course one has to keep in mind that these numbers are estimations published by a group that has Bluetooth's interests at heart. Still, the fact that Bluetooth has become a readily available and widely spread technology can hardly be denied. It is available on the majority of mobile devices shipped today. As was noted in an earlier section Bluetooth operates on the ISM band. While the downsides of this have been discussed, the upside is that it makes the technology available worldwide; because so is the ISM band. I think it is safe to assume that the availability and pervasiveness of the Bluetooth technology is sufficient for the needs of this project.

Infrastructure Dependencies Since Bluetooth connections happen on a point-to-point basis there are no dependencies on external infrastructure. The only thing required for a working link are two Bluetooth-enabled devices. Therefore the technology fulfils all the requirements I have set in this category.

Energy Consumption Most mobile Bluetooth devices are of class three (see table 2.1 if you have forgotten which one that is) due to their need for low energy consumption. As noted in [71] almost all cell phones are of this class. While this radio class indeed allows for extended battery life it also entails a maximum distance of 10 metres. To achieve transmission links that cover the kind of distance this project needs, the use of radio class one would be a necessity. Using up to 100mW of energy, this is another matter entirely. As Sauter states in [87, p. 350] the energy consumption of class one is extremely high as compared to class three; therefore it should only be used where energy consumption is not of critical importance. As for actual real world observations we may look at measurements taken by Lee *et al.* in [62]. They measure the energy consumption of a single chip Bluetooth system to be 57 mW during transmission. Similar values are found by Carroll and Heiser in [37] during an examination of power consumption in smartphones. They quantify the power consumption of the phone's Bluetooth module to be 45 mW at a range of 10 metres. This energy usage is in the same magnitude, albeit a bit higher, as that of ZigBee (see page 35) and considerably lower than those of WiFi (page 28) or GPS (page 31). It is noted in [27] that Bluetooth devices assume an "always on" power profile; they do this in order to be aware of new networks for them to join and thereby reducing the battery lifetime. So in a way the very design of how Bluetooth connects makes it unfit for long operation times without recharge. We may safely assume that a long distance link over Bluetooth is not a viable solution for mobile devices. This will most likely make Bluetooth an ill fit for what is proposed in this paper.

Second Technology: WiFi

An increased need for mobility and flexibility has seen to the rise of wireless local area networks. As of the date of this work the standard IEEE 802.11, more commonly known as *Wi-Fi*, *Wireless LAN* or *Wireless Ethernet*, is the most popular protocol of choice for wireless networks. The IEEE 802.11 standard was conceived in 1990 and accepted in 1997. Its aim is to introduce medium access control and physical layer standards [52]. When speaking of IEEE 802.11 one is actually only just addressing the core and basis of the standard. Since its inception it has grown into a family of standards based on this root. Currently the following 802.11 standards still experience regular usage by end users: 208.11a, 208.11b, 208.11g and 208.11n. All of these

PHY	transmission	range
FHSS	1 Mbit/s	120 m
DSSS	2 Mbit/s	100 m
DSSS	5.5 Mbit/s	55 m
DSSS	11 Mbit/s	40 m

Table 2.2: Typical range of a WLAN AP (unobstructed)

are enhancements on the original standard, introducing improvements to security and performance. But since they also have sometimes profoundly different workings (e.g. modulation schemes) speaking of *the* 802.11 standard is difficult. As was implied above this standard is governed by the Institute of Electrical and Electronic Engineers (IEEE).

Range The typical range of a 802.11 access point station according to [53] is given in table 2.2. The first column specifies which kind of Physical Layer Protocol (PHY) is used, determining how connections are set up and maintained. The throughput and range of the transmission are displayed in columns two and three respectively. Obviously the throughput decreases with increasing range due to worse signal to noise ratios. Note that the measurements in table 2.2 are based on the performance of access points, a type of infrastructure for which energy consumption is typically of no concern; furthermore they only hold for unobstructed paths of transmission.

Further observations regarding the range of DSSS may be found in [51]. Hein writes of possible distances of up to 550 metres on an open field, 115 metres in semi-enclosed environments like a city street and 50 metres within buildings. These are impressive numbers that would surely prove sufficient for my purpose. Sadly the device with which these results were procured is only stated very generically as "Access Point" which makes it difficult to translate these values onto mobile devices. Despite this I see it fit to assume WLAN has the potential to provide the kind of transmission range I am looking for. The energy consumption discussed later will be another matter entirely.

Interference Resilience WiFi operates in the 2.4GHz range, otherwise known as the ISM bands. This it shares with two other technologies introduced in this chapter, Bluetooth and ZigBee, among numerous others. As was already stated, this constitutes a quite hostile environment. Unlike Bluetooth (which utilises frequency hopping) or ZigBee (which may form flexible mesh networks) Wifi does not seem to feature any mitigation technique for this except "when in doubt use more power". Currently there seem to be two main concerns for the quality of a WiFi connection: Bluetooth

The 802.11 Physical Layer Protocols

The 802.11 standard supports several kinds of PHY, two of which are mentioned in this section: *frequency hopping spread spectrum* (FHSS) and *direct sequence spread spectrum* (DSSS). As the name implies FHSS is a technique that periodically changes the broadcasting frequency. It does so in a pre-determined order that both the sender and receiver have to be aware of. The advantage of this method is that a large number of networks may coexist without negatively influencing each other's rate of transmission. As noted in [46] this requirement has become obsolete, ousted by the need for high throughput in each individual network. DSSS on the other hand takes a narrow band signal and spreads it over a wide frequency band. The latter is monitored for changes by the receiver; the original signal is recovered by inverting the spreading process. This approach showed more promise of higher speeds, quickly making it the more favoured PHY.

Side note: the three PHY originally specified for 802.11 also featured one for infra-red. However to the authors knowledge no device that implemented this ever reached the market.

transmissions and microwave oven leakage. Shuaib *et al.* notice in [91] that Bluetooth transmissions may notably influence the throughput of 802.11g networks. For weak connections the reduction in throughput may be as high as 53%. They also find that ZigBee hardly interferes with WiFi transmissions; this may be explained by ZigBees comparably low transmission power. Lakshminarayanan *et al.* note how "*802.11 performance is affected considerably in the presence of Bluetooth*" in [60]. According to [103], microwave oven leakage also seems to be a major source of interference to 802.11 networks. Additionally Gummadi *et al.* found out during experimentations in [48] that a "*relatively small amounts of RF interference from devices that share the ISM band but are not 802.11 compliant can result in substantial performance problems*". WiFi proves to be somewhat susceptible to interferences. This mostly seems to be mitigated through increased transmission power by access points - a strategy that is hardly feasible in a settings with constraints on energy usage and battery lifetime. Nevertheless I have not found any indications to a complete connection outage caused by RF noise. So while WiFi is probably not the most resilient form of communication introduced in this chapter, it should prove sufficient.

Dissemination and Availability WiFi has established itself very well as the solution of choice when it comes to wireless (home) networking. The

WiFi Alliance's homepage claims 450 Million users worldwide [20]. The technology is certainly popular and has become widely available. The IEEE 802 standards themselves are available to the public free of charge six months after their initial publication [7].

Energy Consumption As Lee *et al.* state in [62], "*WiFi is designed for a longer connection and supports devices with a substantial power supply*". More concisely they compare the energy consumption of several devices implementing different wireless protocols. They find that the energy consumption of the WiFi module is larger than that of ZigBee and Bluetooth (undisclosed power class) by a factor of ten. They number the power consumption of the WiFi device at 219 mW during transmission. In [37] Carroll and Heiser examine power consumption in a smartphone. They quantify the power consumption of the WiFi module at roughly 700 mW during a network benchmark. The much higher value may be explained by differences in the hardware implementation of the test devices and the fact that measurements in the latter source were obtained during a network stress test. Nevertheless their value is backed up by Abdesslem *et al.* in [29] where again a mobile phone is taken as point of reference. They number the energy consumption of the WiFi module at an average of 661 mW, a value quite close to that found in [37]. In conclusion WiFi's energy consumption profile is probably not the best fit for the solution this project aims for. The values taken from [37] and [29] are to be taken as upper bounds since they were obtained during constant transmission at maximum capacity. Continuous transmission and high throughput will not be part of my use-cases. IEEE 802.11 is likely not the best choice for my project when it comes to long battery lifetimes.

Infrastructure Dependencies WLAN networks feature two modes of operation: Infrastructure and Ad-Hoc. The first depends on designated access points. With the latter one it is possible to achieve connectivity between two devices without a central managing device. Therefore WLAN fulfils the requirement of being able to establish a connection without a fixed-point or third party infrastructure.

Third Technology: GPS

The *Global Positioning System* (GPS) is a satellite navigation system consisting of 24 satellites on six different orbital plains as well as a ground control and monitoring network. It provides "*accurate, continuous, worldwide, three-dimensional position and velocity information*"[55]. In 1973 satellite navigation programs by the U.S. Navy (the *Timation Program*) and Air Force (Project 621B) were merged into "NAVSTAR Global Positioning System". This new system was supposed to offer three dimensional real time

The beginnings of global positioning

Following the first launches of satellites into earth orbit came the ideas for an extra-terrestrial navigation system. *TRANSIT* was one such system and went into operation for the U.S. Navy in 1964. The main purpose was determination of a ship's position in order to allow it to carry out more accurate missile strikes. TRANSIT utilised the Doppler shift of signals received from the satellites to determine two-dimensional terrestrial positions and featured an accuracy between 200 and 500 metres. It was unable to determine an object's speed or altitude. TRANSIT gave rise to the specifications for the second generation of the U.S. satellite positioning system, which came to be known as *GPS*. [86]

geolocating - as opposed to its predecessor *TRANSIT*. Usage was to be limited exclusively to the U.S. military and such allies thereof as were seen fit. The people in charge of the project were confident that the technologies they utilised were too expensive, the algorithm too complex to be used in the civilian sector. They were so sure about this that no encryption was used on the GPS signal; merely the structure of its data packages was kept secret. As is customary in academia the existence of information of unknown structure paired with an "*you are not able to use this it's simply too complex*"-attitude was seen as a challenge and a bit of rather good sport. Consequently the first GPS receivers for the civilian sector were not far behind. The public usage of the satellite system was later officially sanctioned by U.S. President Reagan. All 24 satellites were in orbit by 1994 and operations were fully commenced the following year. By 2010 as much as 95% of *GPS* usage had fallen to the civilian sector [42].

Range GPS is available worldwide [42][55][86]; therefore range as discussed with other systems in this chapter will not be an issue. But because GPS works fundamentally different than the other technologies discussed in this chapter (i.e. extra terrestrial) accuracy of positioning will be more of a problem. As I have stated in section 2.2 the range requirement is defined as being able to maintain an area of influence within which the technology is able to detect the presence of specific devices. For this a certain accuracy will be required in case of global positioning. Per specifications GPS for civilian use provides accuracy of at least 13 and 22 metres in the horizontal and vertical plane respectively. These values are, however, quite pessimistic as compared to actual performance. Long term accuracy measurements have shown the average errors to be 7.1 metres for the horizontal plain and 11.4 metres for the vertical plain, both values given with 95% confidence. How does this translate to actual mobile device performance? Kaplan and

Hegarty state the median positioning error of a handheld device to be 25 metres on land with a worst value of 54 metres - assuming no shielding and a nominal noise environment [55]. However an urban area may prove a far more challenging testing ground. Even without high-rise buildings and deep urban canyons obstruction poses a problem. In [84] stand-alone GPS positioning is shown to exhibit quite large errors in narrow streets with 5-storey buildings (about 20 metres in height), an environment quite common throughout cities in Central Europe. So what does this indicate for the project? GPS shows quite large magnitudes of error in urban environments, possibly resulting in false positives or late notifications. It has to be assumed that trusting in GPS without such techniques as dead reckoning would lend an overall air of unreliability to my system.

Interference Resilience Since GPS is a satellite navigation system there is a lot of room for all kinds of interferences between sender and receiver. According to [99] those are ionospheric, tropospheric, relativistic and tidal effects as well as clock and multipath errors. However there is sufficient knowledge available on how to mitigate the effects of these interferences. In any event they are negligible because there is a far more problematic kind of interference: physical obstruction. If there is no line of sight from the receiver to the required number of satellites (as is the case e.g. inside a building) then failure of the positioning system is a certainty. As stated in [42][p. 182] this makes GPS an inadequate choice on land if continuous availability is required. Remember that GPS specifications were shaped by military application scenarios on land, sea and in the air. At the time of development nobody predicted the need for indoor positioning. Thus it is simply not achievable with a purely GPS-based approach [86][p. 283]. Even outside of buildings GPS signal reception is an issue in urban environments. In [40] Chao *et al.* found that in Hong Kong, signal blockage poses a severe problem to purely GPS-based positioning. Reportedly, reliable and continuous positioning could not be achieved in dense urban areas. Mok *et al.* state in [70] that "*especially in cities with high-rise building, the satellite visibility is a very critical issue*" - for vehicle navigation in their case, but the point remains. The susceptibility to interferences in the form of physical obstruction is the main drawback of GPS. Frequent system outages are to be expected in dense urban areas with little to no chance of positioning inside buildings. This is not feasible for many use-cases.

Dissemination and Availability In 2006 Kaplan and Hegarty wrote that the civilian GPS had become "*the predominant satellite navigation service in use by millions throughout the world*". They also list several estimations for the growth of the market (and that of GNSS in general). These vary in amounts and timespans, but all of them agree that the market will grow further [55]. Concerning the infrastructure, the NAVSTAR satellites have proven more reliable than was originally anticipated. As a result there was an average of more than 25 operational satellites available at any time between between 1995 and 2006 [55]. Since the U.S. government has declared the GPS a matter of national concern and *Sole Means of Navigation* (within its territory) this is unlikely to change soon. Overall it is very unlikely for GPS or GNSS in general to become obsolete or disused any time in the near future. The amount of reliance put into these systems for navigation alone as much as guarantees that.

Infrastructure Dependencies Of all technologies described in this chapter, *GPS* is the only one that has to rely on external, third party infrastructure. This means the 24 satellites in orbit which are necessary to achieve exact positioning. This extraterrestrial network was maintained by the *Interagency GPS Executive Board* (IGEB) of the United States until 2004. The IGEB was superseded by the *National Executive Committee for Space-Based Positioning, Navigation and Timing* (PNT) which is currently still in charge. According to its website ([13]) the PNT consists of members of several U.S. Departments like Transportation, Defence, Commerce, Homeland Security, State and Interior.

With *Selective Availability* the U.S. government demonstrated that it is very well able to effectively neuter the GPS. Although there is currently no evidence on imminent danger of such means being put into place again, the possibility remains [55][p. 640]. Losing this infrastructure would render the whole project futile as no (or insufficient) positioning of group members could be achieved. Although I do suspect that society would have to deal with more pressing issues than the loss of this project should the GPS satellites ever go dark. There are currently one and a half alternatives to NAVSTAR GPS. The first is *GLONNAS*, a system very similar to GPS developed by the Soviet Union also around the same time. It had issues with availability and precision (which in turn where due to neglect in maintenance) around and after the fall of the Union; partly due to that *GLONNAS* never achieved significant market share in western countries [42]. The other, half alternative is the *GALILEO* system developed by the European Union. It is not a full fledged alternative simply because it was not operational as of the time of this project. However it will be compatible with *GPS* once it goes online [42]. The full constellation of 30 satellites is supposed to be in orbit by 2019 [6]. Considering this I *really* hope to finish my work before Galileo becomes a valid alternative.

Energy Consumption Energy consumption of GPS may be assumed to be quite high. Especially in our use-case where frequent, accurate positioning has to be achieved without relying too much on dead reconing (which

Infrastructure dependence case in point: *Selective Availability*

In 1983 a flight of Korean Airlines was shot down by the Soviet Airforce because it drifted from its intended course and entered restricted airspace. Following this the U.S. government allowed and facilitated the public use of NAVSTAR GPS, thereby providing the means of reliable global navigation. To the U.S. military this posed a problem: They feared that precise civilian GPS receivers could find their way into nuclear weapons outside of the States - this was the height of the Cold War after all. To mitigate this risk artificial errors were introduced into the data sent by NAVSTAR satellites. The result was a greatly increased average error of the GPS for civilian receivers. This measure was called *Selective Availability* (S/A). Just how great was the induced error? Here's an example: earlier in this section I quoted the accuracy of GPS to be 7.1 metres and 11.4 metres for the horizontal and vertical plane respectively. During times of active S/A, these values were in the order of 50 and 70 metres respectively. Despite its name S/A was not selective. It could either be turned on or off; neither the region to be affected nor the amount of error induced could be scaled. S/A went online in 1990 and remained active for more than ten years. Only after continuous rising pressure by civilian GPS users and the U.S. economy was this remedied: S/A was deactivated in the year 2000 and has remained thus since.

may become inaccurate fast). Several tests have been conducted on mobile phones to determine the energy usage of their GPS modules. Since those are mobile and should also have been designed with lowest possible energy consumption in mind, I will use them for exemplary values. In [29] Abdesslem *et al.* examined the GPS power usage on a Nokia N95 smartphone. With active GPS in an outdoor area they measured the phone's energy consumption to be 623mW as opposed to 26mW while all sensors where turned off. Ramos *et al.* state the power draw of an HTC HD2 smartphone to be 387 mW during location sensing with GPS while everything else was in standby [83]. Assuming a battery capacity of 1000mAh that means continuous operation for about 9.5 hours. Carroll and Heiser closely examine power consumption of mobile phones in [37]. They find that the energy consumption of the GPS module alone is about 143mW. Apart from these Savvides *et al.* remarked in [88] on the topic of sensor networks that "*the power consumption of GPS is substantial*", stating that equipping it on sensor nodes is not always feasible when limited power supplies are involved. So GPS draws a fair share of power, although it is possible to continuously utilise GPS for a couple of hours. Presumably this should be enough to get through a field trip of medium length, but no more. GPS certainly cannot gather many points in the battery lifetime compartment.

Fourth Technology: ZigBee

ZigBee is a "*standards-based wireless technology designed to address the unique needs of low-cost, low-power wireless sensor and control networks*" [22]. It is built upon the IEEE standard 802.15.4 and defines communication protocols for low-data-rate wireless networking. The name "ZigBee" references the zigzag dance used by bees to relay messages to their hive; it is supposed to be a metaphor for the way devices find and interact with each other on the network and means to emphasise the protocol's mesh networking capability [44]. ZigBee is a young technology – at least as compared to the others in this chapter. Its first developments go back on project called *HomeRF-Lite* from 1998 which was supposed to result in a cost effective and interference resilient wireless solution. A first draft was done the next year but the project was put down after *HomeRF* itself became obsolete. In 2002 the *ZigBee Alliance* was founded by Phillips, Motorola, Invensys, Honeywell, and Mitsubishi [27] and took up where HomeRF-Lite had left. The stack was split into two parts: the IEEE standard 802.15.4 which provides the physical and medium access control layers and on top of that *ZigBee*, which was to be comprised of the network, security and application layers. The standard was finished on December 14^{th}, 2004 and made available to the public in June 2005 [59].

Range The "typical range" of ZigBee is stated as 100 metres in [44], the same as that of IEEE 802.11b. Similarly, [59] gives the same value for an output power of 0 dBm (recommended by IEEE 802.15.4) and a receiver sensitivity of -92 dBm. This is calculated using Path Loss only so it is at best an upper bound in an open field environment. On their homepage the ZigBee Alliance state transmission distances "*ranging from 10 to 1,600 meters, depending on power output and environmental conditions*' [24]. 1600 metres is a hefty distance and I hope that I will be forgiven for assuming that it will not be kind to battery lifetime. So ZigBee features a range similar to that of WiFi, albeit with a far smaller power footprint (see page 35). It is doubtful whether this range will be sufficient to maintain radio contact between group nodes in a building. But what sets ZigBee apart from other technologies introduced in this chapter is its inherent capability for mesh networking. As noted in [75] ZigBee is "*in the first place intended for mesh networking, where it provides a dynamical self-reconfiguring and self-healing network*" - making this specific network topology one of ZigBee's key features. Safaric and Malaric remark that mesh networking can mitigate the limited abilities of singular nodes by "*leveraging the ability to relay data through nearby cooperating nodes*" in [85]. This is where ZigBee could potentially shine when it comes to the range requirement. Utilising a mesh network it may be possible to mitigate transmission power fall-off inside buildings by relaying signals through group nodes to the group leader. It has to be noted that such a thing could indeed be also achieved with Bluetooth, WiFi or UWB, but the ZigBee standard makes it potentially much easier. This possibility will certainly be investigated further during the course of my work.

Interference Resilience As you may have expected if you have read the sections about WiFi and Bluetooth, ZigBee suffers from similar problems as these when it comes to RF interferences. As an added difficulty, ZigBee features much lower transmission power than WiFi (or Bluetooth, potentially). Because of this, when the three technologies coexist in a physical space WiFi, is the dominant source of interference, as was shown in [90]. In [91] the throughput of ZigBee is stated to suffer a ten percent reduction if the channel of operation collides with that of a nearby WiFi network. Worse for my use-case, Lau *et al.* indicate in [61] that localisation systems based on ZigBee are susceptible to WiFi interference. In their experiments with such a system, which utilises the RSSI of ZigBee transmissions, they find that "*even under an ordinary amount of background WiFi traffics, the localisation system might suddenly perform poorly*". They provide evidence that RF interference poses a significant problem to localisation accuracy. This is due to the loss of beacon messages as result of collision with other traffic in the air. While this all does not bode terribly well for the interference resilience of ZigBee, the motto may be "mesh networking to the rescue" here once more. After investigating the interoperability of WiFi and ZigBee for home automation systems in [47] Gill *et al.* summarise that in an environment where "*interference is constantly fluctuating, the advantage of increased communication routes available through the adoption of a mesh topology outweighs the added routing complexity*". They also practically prove that the co-existence of WiFi and ZigBee is possible. On a similar note, Baker states in [27] how ZigBee has advantage due to the ability of mesh multi-path transmissions and hybrid network configuration options. To sum it up, ZigBee seems susceptible to RF interferences by other devices broadcasting on the ISM bands. This is not helped by the low transmission power. But where Bluetooth has frequency hopping and WLAN simply uses more power, ZigBee could utilise mesh networking to overcome this problem through redundancy by multi-path transmissions. Still it has to be noted that this is probably the weak point of ZigBee as far as my use-cases are concerned.

Dissemination and Availability The overall outlook for the development of ZigBee seems to be optimistic. In 2007 Baronti *et al.* predicted that it would probably be "*embedded in a wide range of products and applications across consumer, commercial, industrial and government markets worldwide*" [28]. A similar view is shared by Guo *et al.* in [49] where they state their opinion that "*ZigBee has a bright future*" and how ZigBee wireless sensor networks will become dominant. Indeed a lot of research seems to have been conducted over the past few years concerning the ability and possibilities of ZigBee, giving it something of an "up-and-coming" air. On its homepage [23] the ZigBee Alliance projects the annual shipments of ZigBee devices to reach 350 million by 2016. They announce to have reached the 600 mark of certified products in October 2012. Based on the current developments and projections it does not seem likely that ZigBee will disappear soon. While the technology is not as established or widespread as Bluetooth, WiFi or GPS I see no reason to fear it becoming obsolete or disused. The reasoning behind my assumption is that ZigBee occupies a niche which it seems to fill quite well: very low-energy, accessible and affordable sensors capable of flexible and robust networking. With that it sets itself apart from other existing wireless solutions and could very well find a market.

Infrastructure Dependencies A ZigBee network has no dependencies on any external infrastructure. This it has in common with Bluetooth, WiFi and UWB. Therefore like the aforementioned it fulfils the requirements I set concerning this point.

Energy Consumption Energy consumption is one of the strong points of ZigBee; its "stronghold", if you will. Consequently it exhibits the least energy consumption and longest battery lifetime of all the technologies discussed in this chapter. The ZigBee radio used by Lee *et al.* in their comparative study of wireless protocols exhibited current consumption of 25 mA and 27 mA while transmitting and receiving, respectively [62]. They find WiFi to use almost ten times and Bluetooth still almost double the current consumption of ZigBee in their report. Guo *et al.* describe similar findings in [49] where they state the current consumption of ZigBee to be 30 mA during transmission and less than 0.1μA during idle times. They too state that WiFi and Bluetooth use more power - much more in case of the former. To sum it up this is one of (if not *the*) strong point of ZigBee and it certainly fulfils this requirement better than anything else described in this assessment.

Fifth Technology: Ultra Wideband

Ultra Wideband is, as the name implies, a technology for transmitting data over a large frequency bandwidth. Per definition any transmission scheme featuring a signal bandwidth of more than 500 MHz or a fractional bandwidth of 20 percent is to be considered ultra wideband [65], [100]. The term was coined around 1989 by the U.S. Department of Defence. By that time advances in this field were either largely funded by the U.S. government or undertaken directly by the military; as a result much of the development was classified. This changed around the late 1990s when UWB technology started experiencing commercial use and a more rapid development. A milestone came in 2002 with the publication of FCC (U.S. Federal Court Commission) rules which contained radiation limitations for UWB transmissions and permitted utilisation of these on an unlicensed basis [65]. Europe allowed the use of UWB under slightly different and more restrictive conditions (mostly concerning the available bandwidth) in 2005 [74]. Ultra wideband has some promising traits when it comes to signal propagation, especially indoors (more on that later). Apart from those it also enables low complexity, low energy and low cost components through its signal transmission technique: the transmitter broadcasts on a wideband in very short pulses. Due to this, the signal propagates well without the need for conversion to carrier frequencies or amplification of the signal (as is common with other radio transmitters). Likewise this means that the receiver does not need to take any actions to revert those changes. As a consequence it is possible to design simpler components [76].

Range UWB exhibits a couple of characteristics that are favourable when it comes to indoor environments due to its operation at baseband level and available frequency bands. Most notably are multipath immunity [66] and a better capability to penetrate obstacles [100]. Therefore it is expected that ultra wideband would perform better in a closed environment with obstacles than the other protocols discussed in this chapter. That being said, actual performance measurements were not so easy to come by. The nominal range of UWB is stated as ten metres in both [62] and [82]. This is also the range proposed in the original FCC rulings for a throughput of 110 Mb/s. Luo and Law tested an indoor positioning system under the influence of heavy multipath and achieved an error of under half a square metre at a range of 20 metres [66]. However according to Jingjing *et al.* it is entirely possible to use ultra wideband for long-range communication without violating regulations [98]. They describe an impulse radio transmitter that would enable wireless sensor networks with a range of 200 metres. Apart from that there also exist commercial solutions that feature maximum distances up to and above 100 metres, e.g. [14], [21]. All this indicates that ultra wideband could very well fulfil the range requirements for this project. I would even go so far as to call it the best solution where pure achievable range is concerned.

Interference Resilience As was already mentioned, UWB exhibits good material penetration capability due to the nature of its transmissions [89]. This means immunity to passive interference like rain, fog, clutter or metallised strips and the like [45]. Therefore resilience to interference by physical, environmental obstacles is expected to be better than with any other transmission technique examined in this chapter. UWB is also the only protocol (apart from GPS, of course) which is not bound to the crowded 2.4 GHz ISM band in Europe. This has both positive and negative aspects. The upside is of course that the signals are not bound to a frequency band that is already home to quite a lot of other transmission protocols, thereby evading some interferences. The downside is that because of the wide frequency band utilised by UWB there is a good chance it might overlap with several frequencies used as carriers by other technologies; this opens it up to interference from multiple sources. It seems, however, that ultra wideband is somewhat resilient to interference from both wide- and narrow-band sources [101], [102]. Another favourable feature of UWB is its immunity to multipath fading and cancellation [45], [98]. Win *et al.* have tested this in a real-world scenario in 14 different rooms and hallways [97]. Once again this is to be attributed to the short-pulse and wide baseband nature of the transmissions. As an added bonus UWB may exhibit low interference to existing or legacy systems if configured properly [45]. It would seem that ultra wideband offers a high resilience to interference both by other radio transmissions and harsh environments. Its long use for military applications also supports this assumption.

Dissemination and Availability Ultra wideband itself is not a novelty. In fact it has been around since the 1960s [65], [100]. But its most recent incarnation, due to renewed interest, began when the FCC laid down the rules for UWB in 2002. The ECC of the EU has decided on similar rules. The emphasis here is on the word "similar". According to [67] and [74] Europe's approach is more cautious and restrictive than that of the United States. Most notably it allows for a smaller bandwidth of frequencies to be used: 6 – 8.5 GHz as opposed to 3.1 – 10.6 GHz. Japan and Korea also have their own set of restrictions on emission levels on certain frequencies. While it is common for radio transmission regulations to be varying greatly in different regions the problem with UWB is that in many cases those do not seem to be decided yet. It seems like UWB is still being measured up by many official bodies and no international agreement has been reached yet. This is bound to slow market penetration of UWB down. Commercial availability and standardisation is the weakness of UWB at this point in time. While it has gained much academic interest over the past decade, UWB is not yet represented well on the consumer market. As was already mentioned there are companies doing industrial solutions based on UWB (e.g. [21], [14]), but to my knowledge there is currently no UWB communication or sensor system available to consumers. At the point of this writing UWB exhibits very promising characteristics but its future concerning standardisation and commercial usage seems unclear.

Infrastructure Dependencies Like Bluetooth, WiFi or ZigBee UWB is not dependant on external infrastructure to conduct communication between devices. Therefore it fulfils the requirements of this point.

Energy Consumption The short-pulse transmissions of UWB have the advantage of low duty cycles. This means there are only very short bursts of high energy output followed by long phases of idle time, which is good for battery lifetime [45]. This is also helped by the baseband nature of UWBs transmissions which allows for omittance of the signal conversion to carrier frequencies. Thus ultra wideband is expected to have very low power consumption [65], [100].

Technology Assessment Conclusion

In this section I have introduced five techniques usable in wireless sensor networks. All of these have different areas of designated usage. The aim was to assess each of these against a set of predefined requirements so as to find the one best fit for the application scenarios described in Section 1.1. In conclusion to this assessment, I have decided to use *ZigBee* as the technology to realise this project. The reason for this choice are as follows:

- *Low power consumption*: ZigBee transceivers exhibit very low energy consumption. They were designed, among other use-cases, for wireless sensor networks and as such feature very long battery lifetimes. This is an important factor for longer trips without the possibility of recharging.

- *Low cost*: ZigBee devices targeted at wireless sensor networks are usually designed with scalability in mind. As a result there are affordable IEEE 802.15.4 and ZigBee transceivers available. For example, a current generation Arduino ZigBee (which will indeed later be used in my prototypes) is priced at $25.95 [16].

- *Mesh networking capabilities*: The ZigBee protocol has an inherent support for mesh networking. While this would indeed be possible with other technologies discussed in this section, ZigBee offers a more versatile networking capability. With this it should be possible to overcome weaknesses such as low range within buildings.

The closest contender for technology of choice was *Ultra Wideband*. It promises low energy consumption paired with robust radio transmissions and a greater maximum range than ZigBee. So by all accounts ultra wideband transmissions should be superior to ZigBee. The reason why I decided not to use it was a lack of availability on the consumer market. ZigBee is based on two standards, both of which are maintained by widely known and supposedly well funded entities. This should make it easy to design a network of sensors. On the other hand there is no standardised transmission protocol based on ultra wideband. Or at least there was none at the point of this writing that the author was aware of, even after some research into the topic. Indeed several regulatory bodies around the world cannot even seem to agree on a set of rules regarding UWB. If it were not for this, if there was a standardised communication protocol based on this technology, then ultra wideband would have been my first choice hands down. In fact it pains me a bit that this project is conducted maybe a few years before UWB makes a commercial breakthrough. It would certainly enable a lot in the way of interference resistant sensor networking through short-pulse wideband transmissions. Additionally it would have offered high resolution positioning, a feature that would certainly come in handy. I dare say that ultra wideband communication will be very interesting for future improvements to endeavours like this. In the end it comes down to the fact that ZigBee currently seems much more convenient.

Apart from ZigBee and UWB, *Bluetooth* would have been a possibility to work with, albeit far from optimal. The intended application profile seems to diverge too much from what I had in mind to be a straightforward technology of choice. As I have said, Bluetooth was designed to replace physical cables, something it does well for e.g. hands-free mobile phone equipment or computer peripherals. Because of this it aims for a high throughput and only short range in most cases. In contrast I will most likely need very little data transmitted but as much range as I can get without sacrificing battery life. Additional Bluetooth is more complicated than any of the other three terrestrial protocols in this chapter [62]. It probably would have been possible to make do with Bluetooth, but most likely also much more complicated.

Next in line would have been *WiFi*. It suffers from similar drawbacks as Bluetooth in that it was not designed with low power sensor networks in mind. Instead it was intended as replacement for physical cable infrastructure and personal area networks. As such it focuses on high data throughput, much more than is needed for this project. The energy consumption would be too high to make sense for my scenarios. Apart from this the hardware would still be more expensive and probably also more cumbersome than that of ZigBee, UWB or even Bluetooth.

Lastly there would be *GPS*, or GNSS in general. The biggest drawback here was that satellite navigation will not work when line of sight to the required number of satellites is lost. This means no functionality within buildings. There have been suggestions on how to mitigate this by ways of dead reckoning or combining GPS with WiFi for indoor localisation. Some of these approaches are described in [86], for example. All of them have in common that they combine GPS with other technologies, thereby making the whole system more complex and adding more potential points of failure. In addition to that, GPS also exhibits relatively high energy consumption, reducing battery lifetime to a few hours. At the most, GPS could be considered an additional sensor to use while the system is operating outdoors.

2.3 On Estimating Distance and Position

With the rise of wireless sensor networks, the topic of estimating the position of a node within such systems has received much attention. There exist various ways to try and locate the source of a radio signal. In this section I will provide a short description of the basic theoretical concepts of localisation and range estimation; furthermore I will also give an introduction in to the current state of research on these topics.

Let us begin by establishing the basic theoretical part of range estimation, since it is the foundation for localisation. Several techniques are known for determining the distance between two nodes based on the radio transmissions between them. They can be roughly categorised into two groups: those that rely on the propagation time of the signal and those that are based on the signal's quality on the receiving end.

The first method found in the former category is *Time of Arrival*, shortened as ToA [68]: One node sends a signal to another node and the distance between the two is derived from the time it took the transmission to arrive – also known as time of fligh. The thought is simple, the practice is less so. For ToA to work, both nodes need to have quite accurately synchronised clocks. This demands both a synchronisation mechanism within the network and hardware support for very precise clocks. Both things will increase the cost and complexity of the sensor network [68]. A variant of the ToA method is called *Round-trip Time* (RTT), which changes the principle to a two-way propagation model. Instead of calculating the time of flight, the receiving node simply bounces a transmission back to the original sender. The latter then tries to derive the distance from the time difference between sending the transmission and receiving a response, i.e. the time it took the transmission to do a round-trip (hence the name). This rids the system of the necessity to synchronise all clocks, but in exchange also adds an uncertainty factor: the time it takes the receiver to bounce back a message to the sender (i.e. the time between receiving and sending again) may be unknown and vary. A third method besides ToA and RTT is for one device to simultaneously emit two signals with different propagation characteristics, like for example a radio transmission and some form of sound. The receiver can then derive a distance estimation based on the timespan between receiving the "fast" and the "slow" transmission. This form of distance estimation is known as *Time Difference of Arrival*, or TDoA. The problems arising from TDoA are that a) the sender and receiver need to be able to process different types of signals and b) that the range of one of the signals is usually quite low, between three and ten metres [33]. What all of these time-based systems have in common is the need for hardware that is able to measure very small time differences, due to radio signals propagating at the speed of light.

A different approach to estimating distances in wireless networks is measuring the signal quality of incoming transmissions. The strength of a radio signal will decrease as the distance between sender and receiver increases; ideally, the former is inversely proportional to the latter, squared. Thus, if the output power at the source of the transmission is known, one could infer the distance between the two points by measuring the strength of the signal at the receiver. However, this measurement, called *Received Signal Strength Indicator*, or RSSI, is highly susceptible to distortion by obstacles. It is noted in [31, p. 4] that, "*in practice, environment influences such as unpredictable reflections, interferences and obstacles as well as unknown electronic transfer elements within the transceivers impact the measured data enormously. Hence, calculating distances using simple models will result in faulty distances with at least offset errors and scaling errors*". Two different approaches to smooth out the RSSI data are explored in [57]; both help to improve the correlation between signal strength and distance. Still, basing distance estimation on the signal strength provides less accurate results than using time differences. The big advantage is of course that RSSI readings are easy to obtain, thereby helping to save costs in both development and hardware; hence this technique is popular with many wireless network projects, and still remains promising according to [38]. A variant of of RSSI is the *Packet Loss Ratio*, which tries to gauge the signal quality by counting, as the name implies, the packets lost in transmission. This technique is evaluated in [104], where it was found that it suffers from the same inherent unreliability as RSSI readings, exhibiting variations in different directions (not distances) from the sender and being susceptible to interferences. RSSI and Packet Loss Ratio are sometimes combined for a *Link Quality Indicator*, though I have seen the term used as synonym for signal strength alone also.

Now that we have established the methods for gauging distances it is time to take a look at the localisation techniques, most of which use one of the approaches listed above as a basis; several of these techniques are known and have been discussed in literature to a great extent. Be aware that they are discussed here for the sake of completeness only; localisation will not be part of my work any further, for reasons explained at the end of this section. The most basic and well known localisation technique is probably *triangulation*, whereby a node's position is estimated by calculating the intersection of three circles. For this to work, we need to derive the distance of the node to three reference points of known position (called *anchors* or sometimes *beacons*); the means for this were discussed in the previous paragraph. We may then simulate a circle around each of the anchors, with the distance between it and node being the circle's radius. The point where all three circles intersect represents the position of the node. As is stated in [33], the problem with this approach is that one single, clean intersection between the circles exists only in theory. In a real world scenario, the inaccuracies inferred by the range estimations cause an infinite number of possible solutions when trying to determine a node's position. Apart from triangulation there is also *trilateration*, where a node's angle to at least three anchors is determined and then used to calculate its position. This is actually the principle that GPS works on, with the satellites being the anchors and the receiver being the node of unknown position.

A lot of algorithms have been discussed in literature that aim to determine a device's position in a wireless sensor network. Most of them aim to keep the number of anchors low, as these tend to increase the network's cost through additional hardware (e.g. GPS modules). For example, [80] gives a good introduction into (and also a performance an evaluation of) the topic of *cooperative localisation* in wireless sensor networks, whereby measurements between any two nodes, whether their position be known or not, are used to aid in the localisation. A localisation method called *recursive position estimation* is introduced in [26], where there are only few anchors, but each node that derives its position from them becomes a reference point itself, thus significantly extending the range of the network. An approach that only uses one single yet *mobile beacon* is described in [92]; this beacon needs to traverse the entire network while broadcasting its position, thereby enabling nodes that are in range to estimate their own positions so long as they receive at least three different updates from the anchor. A different form of localisation is proposed in [105], where a self-organising algorithm draws up a *relative coordinate system* containing the network's nodes without any centralised knowledge of the topology; "relative" in this case means that there is no way to transfer the network's coordinate system to a geographic one unless some form of GPS positioning is used on several nodes.

The list of different localisation algorithms discussed above is by no means exhaustive; it merely contains those approaches which I stumbled upon during my research and deemed interesting. Despite the many possible ways that have been described to tackle the problem of determining a node's position, there is none that can be used in my project. Even if I had the hardware to perform distance measurements accurate enough, all of the algorithms for localisation (apart from [105]) require something I cannot provide with this prototype: anchors of known position. Thus the endeavour of locating the child devices with nothing but the ZigBee network is rendered moot from the beginning.

Chapter 3

Assembling a Prototype

Please note that I will not cite actual program code in this work, as it does not contain any great insights or research value. The corresponding sections will outline the control flow of my programs and explain some solutions, but I see no point in throwing excerpts at you that make up a very small part of the actual program and lack all context. Program code should not be the main focus of this work.

After having evaluated my options for a wireless technology and gathered feature requirements for the device in the previous chapter, I will describe here my work on assembling a prototype system. This chapter will illustrate the technical details, both in hardware and in software, and document the user interface in its various stages throughout the development process. If this chapter sometimes reads like a development diary, than that's because it is one, to some extent. I did choose to write it in such a fashion as an introduction to hardware prototyping for those readers that, like me, do not have much experience in these matters.

3.1 Main Device Prototype

I started my work on the first functional prototype as soon as the requirement analysis interviews and technology assessment were done.

Choice of Hardware

At the very beginning, I had to make a choice on how to develop the prototype. Since I am not apt in electrical engineering, neither through my studies nor through personal interest, doing low-level implementations like designing the basic circuit boards myself were not viable. Therefore I decided to go with a ready-made toolkit. A comparison of existing prototyping systems may be found in [41]. My choice fell on *Arduino*, which is an open source project available under the *Creative Commons Attribution Share Alike 3.0* license [10]. The name *Arduino* refers to both the hardware and the development environment used to program it. The reason for my choice was mainly the popularity of the toolkit at the time of this writing. It had already gained a considerable user base, which meant that there was a lot of know-how available through the community - something that should never be underestimated when one is trying to learn a new tool.

Excursion: Origin of Arduino

The Arduino Project was initiated as an educational platform in 2005. Originally it was created for a class project at the Interaction Design Institute Ivrea [25]. The basis for Arduino was *Wiring*, another open source board with various connection capabilities. Wiring already provided a user-friendly development environment (using the *Processing* language for programming) and a ready-to-use circuit board. The Inception of Arduino came from the increasingly demanding needs of students and teachers. The platform of choice up to that point had been, for about a decade, the *BASIC Stamp*. But its microcontroller recently was unable to fulfil the processing needs of new projects and the price tag of about 100$US was not a small one. Thus the Arduino project was initiated by an associate professor and some of his peers to provide students with a better yet cheaper alternative. In the beginning, Arduino spread mostly through word of mouth in academic circles. Slowly the idea of prototyping for everybody took hold, though. In 2010, more than 100,000 produced Arduino boards were reported [30]. In late 2011, a different source already tells of 250,000 shipped units [25]. Arduino seems to have found a firm foothold with the do-it-yourself crowds of various disciplines.

Its open source philosophy (both in hardware and software) also encouraged the availability of many third party peripheral components (called *shields*) to be available in addition to the already sizeable array of official Arduino hardware. A quick search on the Internet revealed that there were multiple providers for all of the functionality I needed. Furthermore, Arduino features a brand of components intended for integration into wearable items, the *Arduino LilyPad*, which I expected to best serve my needs when implementing the child devices. More on my work with the LilyPad may be found in Section 3.2.

After having found the right prototyping toolkit, the next step was to compile a list of components - a "shopping list", if you will. This was indeed not as trivial a task as it may sound. I did not have any prior experience with Arduino prototyping, yet I needed to order quite a few components for a complex system. As an added difficulty I realised that, to get all of the parts I wanted, I would have to order from a vendor in overseas. Especially the LilyPad components for interfacing with XBee modules were hard to come by in Europe.

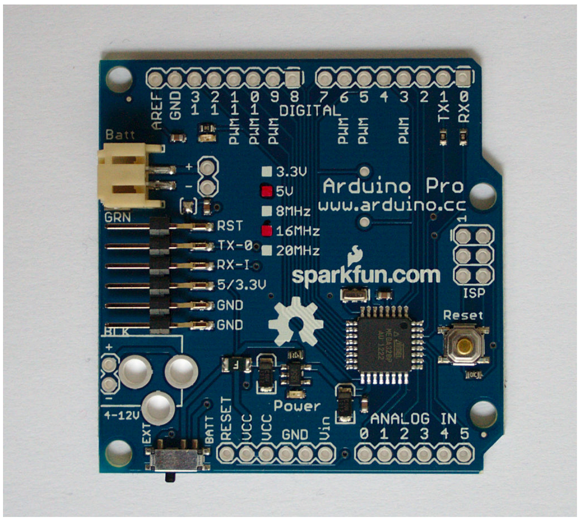

Figure 3.1: The Arduino Pro board as used for the main device prototype

There were many different flavours of Arduino microprocessor boards available; the official homepage featured no less than nineteen in July of 2013. They differ in the built in microprocessor, the form and design of the circuit board, the types of connectors for programming and attaching components as well as the operating voltage (both input and output). For my project I used and *Arduino Pro* board featuring an Atmel Mega 328 microprocessor with a speed of 16 MHz, two kilobytes of SRAM and 32 kilobytes of flash memory running on a 5V input. The item is shown on Figure 3.1. The reasons for this choice were the small form factor of the board as well as the lack of pre-soldered headers, meaning that I could save further vertical space.

For a display I chose a *Nokia 6100 LCD knock-off* fitted on an Arduino compatible shield. The display provides a resolution of 128x128 pixels in 4,096 colours. It runs on the standard 5V output provided by most Arduino boards (the aforementioned Arduino Pro board is among them, of course) and comes with an integrated LED backlight. Since this is cellphone technology I hoped that it would provide good readability in most lighting conditions - I was not disappointed. You can see it in Figure 3.2. The header-less design again provides more flexibility and possibly saves vertical space when connecting the components. Sadly, I did misjudge the size of the display somewhat; the screen itself is only 30.5 millimetres high and wide. It is very well readable, but I was surprised by how small it

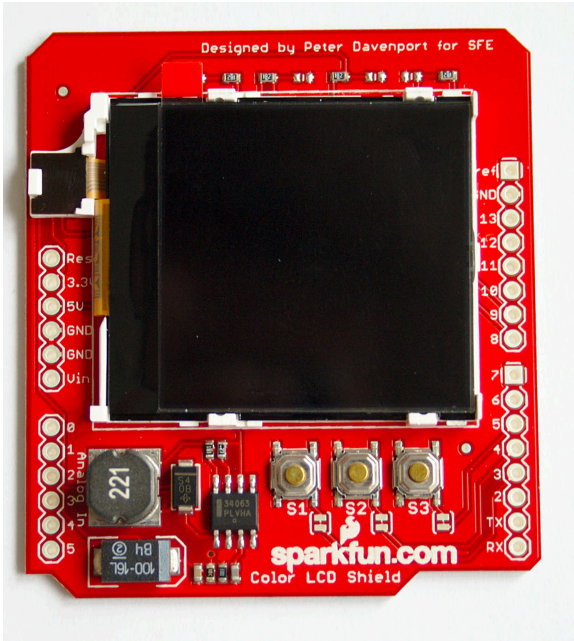

Figure 3.2: LCD screen mounted on a shield

was[1]. Ironically enough, the Arduino project released an official LCD shield in Spring 2013, about three months after I had purchased all components. This screen was bigger, had a better resolution and featured an SD-Card slot. To top that all off, it is also cheaper.

The third component to be fitted into the main device body was the shield hosting the radio module. Arduino provides its own implementation of the ZigBee protocol, called *XBee*. For the XBee radio modules themselves there was only one provider that I was aware of at the time of the project. I chose the Digi XBee Series 2 modules with built-in antennas. These are fully compatible with the ZigBee standard and thus also provide the same mesh networking capabilities. They provide about 2mW of output power. One such module is depicted in Figure 3.23 on page 67. In hindsight it would have been more prudent to purchase the variant with wire antennas to achieve better range. To mount the module I used a shield by SparkFun (Figure 3.3) that seems similar in layout to the official Arduino wireless proto shield.

Last of the big parts of the main device was the battery. I opted for a Polymer Lithium Ion based rechargeable providing 3.7V at 2000mAh. I was sceptical that this current would be able to power the whole device (it did, in the end), but the item seemed a popular energy source for Arduino

[1] Admittedly, that is mostly due to me not properly reading the vendor's specifications rather than some negligence on their part.

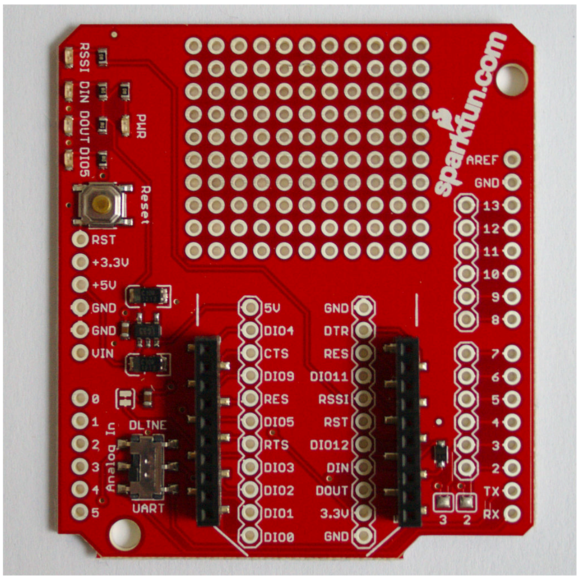

Figure 3.3: XBee shield

projects, so I decided to go with it. The battery was very slim, taking up little space and weighing only 36 grams. Smaller parts on the list of things to be fitted into the prototype casing were an electro-magnetic transducer, one double-throw switch for closing or breaking the power circuit as well as three pushbuttons for navigating the user interface.

After having received all the parts, I went out in search of a chassis to hold them. Since most of the ready made casings for Arduinos do not seem intended for mobile applications, this search came down to traipsing around hardware stores in search for a best fit. After some searching through several general purpose enclosures I did indeed find a sturdy yet comparatively lightweight hard plastic case that fitted the width of the Arduino parts very well. Figure 3.4 shows the chassis as purchased, without modifications.

Assembly

Since I now had all of the parts I needed in hand, it was time to start with the assembly; the first thing to be addressed were the device's innards. There was little vertical space to be had within the chassis, since it needed to hold the Arduino board together with two stacked shields. The chassis was just short of 3cm high; with the enclosure being 2mm thick, this did not leave that much space to go around. Luckily I had ordered the parts without headers being pre-soldered onto them, considering that the typical[2] Arduino headers for shields and boards feature a height of just under 2cm. I opted for

[2]Insofar as these or similar often come pre-soldered on a lot of Arduino items. Many of the official Arduino boards and shields have them in place already at the time of sale.

Figure 3.4: Original casing as bought from store

a different approach and purchased a row of breakaway headers that were just 11.5mm long. The short end I soldered onto the main board, while attaching the LCD shield to the long one. Thus the screen was mounted onto the the Arduino with only a minimal distance between them; just enough to avoid contact.

The XBee shield was another matter. I did not want to attach it between the main board and the LCD shield, because placing a radio transmitter right in the centre of a device will not do any wonders for the signal strength (of which there's fairly little to start with). There was also the option of placing it beneath the main board with the XBee transmitter facing upwards; yet the mount for the wireless module is quite high. Therefore connecting it to the underside of the main board via headers would taken too much space, and also again place the transmitter module between two other parts and more at the centre of the device. In the end I flipped the shield, placing the XBee module against the device's wall, and connected the required outputs with the main board through wires. The end product of this step is shown in Figure 3.5. The three wires protruding from the bottom will later be connected to the buttons. The two longer wires running to the bottom off the screen would be connected to the buzzer.

The next step was adapting the chassis to accommodate the electronics. This mainly consisted of creating an opening for the LCD screen as well as the three pushbuttons and the power switch. I originally intended to also

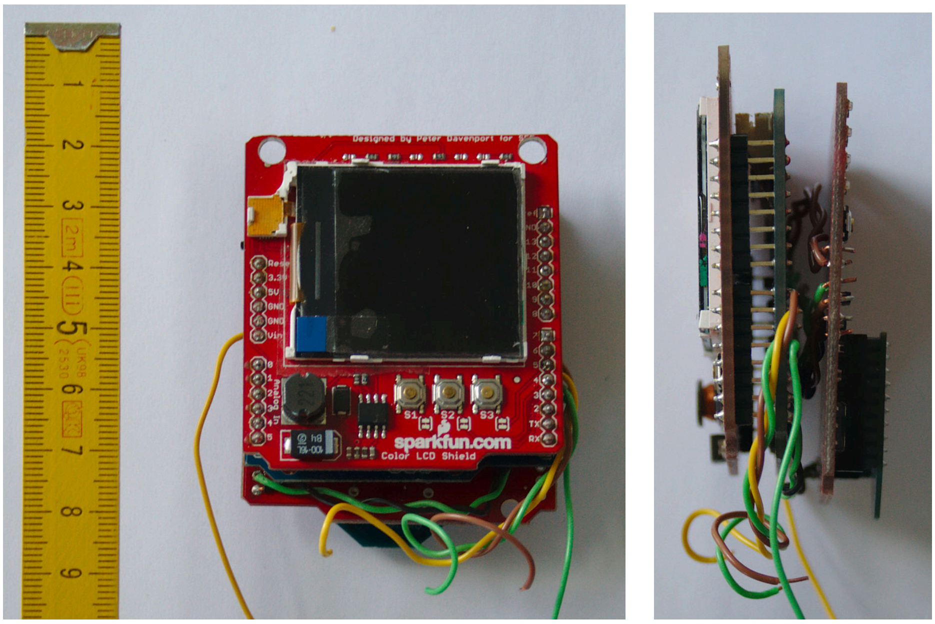

Figure 3.5: Electronic innards of the main device

include a bay for the XBee module to protrude from the device's backside. Ultimately the decision against that came from considerations for splash water resistance and the possibility that the module could fall off and be lost (it is not soldered onto the shield, after all). You can see the adapted casing in Figure 3.6. As you may notice, there is an additional hole near the bottom left corner of the screen recess. This was necessitated by the fact that there is a rather large inductor on the LCD shield that is higher than the actual display (may be seen on Figures 3.2 and 3.5) - a rather inconvenient construction method for practical purposes. The three screws that can be seen in the left image (depicting the inside of the casing) keep the electronics in place. You may have noticed that two holes were drilled into the LCD screen when viewing Figure 3.5. These fit with the two screws near the upper end of the device. The third screw beneath the display recess stabilises everything by "wedging" the components in.

The remaining task of installing the battery, power and buttons was very straightforward. Figure 3.7 shows the device's innards with and without the battery pack (3.7b and 3.7a, respectively). The power switch is at the bottom of the device and is better visible in Figure 3.14 on page 59. That concludes this part of the hardware description. The whole exterior of the prototype as it was used during the usability tests can be seen in Figure 3.8.

A word on the use of hardware buttons: I am well aware that touch-based interfaces have gained incredible traction, mainly fuelled through the

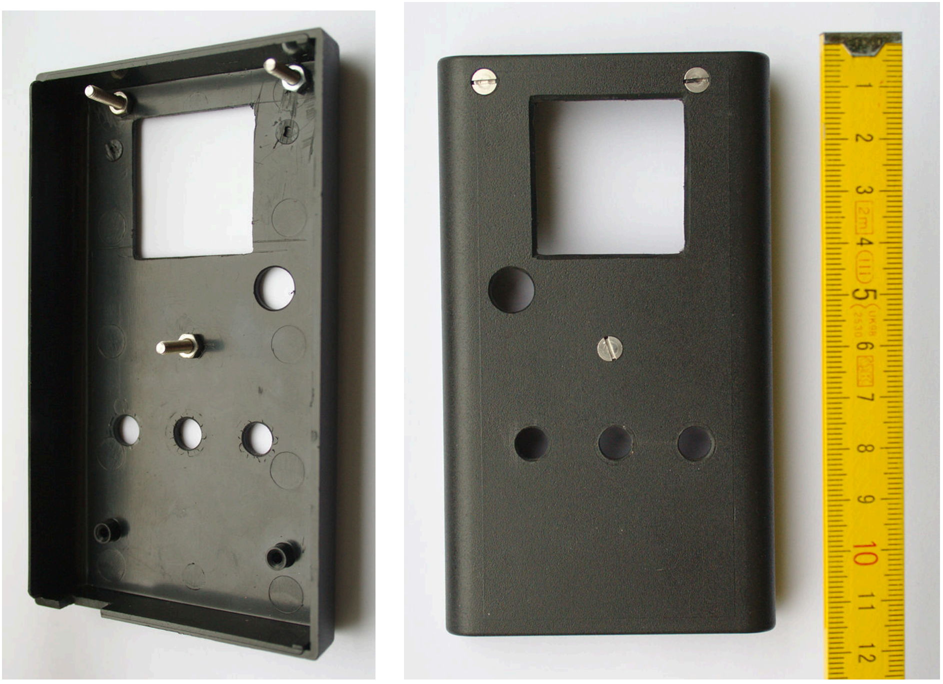

Figure 3.6: Casing after modifications

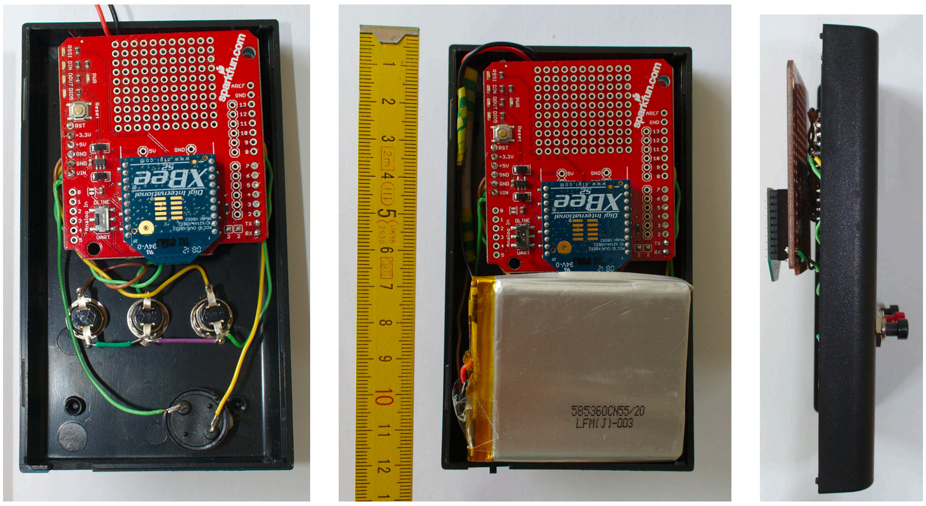

(a) no power circuitry (b) including battery (c) side

Figure 3.7: Arduino electronics fitted into the casing

rise of smartphones in recent years. I have also thought about using this input method with my prototype; the required hardware is available for Arduino, and there are some software libraries floating around the internet which facilitate addressing and programming it. However, I opted to

Figure 3.8: Exterior of the main device prior to usability testing

go with physical buttons, for the following reasons: they provide a tactile handling and feedback that enables fast and sure usage even without visual contact. Touch-based input methods, on the other hand, require the user's undivided visual attention. As an example, take alarm clocks as opposed to smartphone alarms: you can't easily end a smartphone's alarm without looking at the display; yet with a "traditional" alarm clock, featuring a dedicated hardware button, you could do that interaction blindly. Another advantage that physical buttons have over touch-based input methods is their resilience to rough environments. A touchscreen that is wet or dirty quickly becomes unusable. During cold weather a touchscreen is difficult if not impossible to use with a gloved hand. Physical buttons are just much more resilient to these adverse conditions. They work regardless of how dirty the device's surface becomes. In the end, it comes down to this: in a hectic situation, where the user is already performing other tasks (like, say, watching over a group of school children) and has a high cognitive load to deal with, a purely touch based input system only adds to this problem [34] and in addition is potentially more vulnerable to dirt and grime.

Software

For programming the main device, I stuck to the standard Arduino IDE, mostly for its ease of compiling and uploading code. I used three libraries in addition to the core functionality: *XBee* [4] is a library that adds a layer of abstraction and convenience to working with XBee modules. It facilitates an easier way of transmitting and receiving data over wireless as well as programming or reading parameters of the XBee module during runtime. Most notably, it takes over the burden of having to assemble valid ZigBee messages one byte at a time. *StandardCpluspus* [17] aims to bring all the features of `C++` to the Arduino environment, as implied by the name. The reason for including this is that it provides the concept of maps for storing pairs of values. This was my preferred way of managing the group member list, reporting times and signal strengths. That would, of course, have been possible through simple arrays, yet much more cumbersome. Lastly, *gLCD* [3] provides some functions to easily program the LCD display I am using for the prototype. It features rendering of simple geometric figures (lines, rectangles and circles) as well as text in two different sizes and letter spacings. Additionally I used the *SoftwareSerial* library [15], albeit for debugging purposes only (at least on the main device). It allows to output a serial interface on two arbitrary Arduino pins. This is invaluable for debugging since the XBee module already hogs the "regular" serial[3]. Everything else was core functionality provided by the Arduino environment.

The general flow of the program is quite straightforward. The `setup()` function initialises the XBee library on the serial interface, defines the three pins used by the buttons and starts the interrupt timer for reading them (I explain about interrupt timers later in this section). It also initialises the LCD library with the right display type. Lastly it makes the buzzer sound once to indicate a ready state. The `loop()` method first listens for incoming transmissions with a timeout of 200ms. If it receives a packet during that time, it reads the request type indicator (set by the sender in the payload of the transmission) and acts accordingly. There were only three request types implemented for this prototype. Please refer to Section 3.3 for more information on the network infrastructure. After potentially responding to requests, the program will check for absentees, provided that a certain time has passed since that was done last. If anyone is overdue, meaning no report was received for more than a given number of seconds, it will consider that member lost and sound an alarm. Afterwards the routine checks when the user has last made an input, i.e. pressed one of the three buttons. If a certain time has passed between then and now, the device will go into a locked mode. The display is turned off (safe for the backlight, which seems to be always on) until the user presses a button again.

Since I already mentioned it, on the topic of button presses I have the following remarks: Arduino does not support any form of concurrency, so there is no way of setting up a dedicated listener job that watches for button presses. Checking the buttons in the `loop()` function is also a bad idea, since that may take a variable time to complete, thus leaving the system seemingly unresponsive to user input if one loop is taking a while. What the system

[3]Most Arduino devices have a dedicated serial interface, on pins zero and one, which is used to communicate with various peripherals. This native serial support happens via a Universal Asynchronous Receiver/Transmitter (UART), which allows for serial communication even while the microprocessor is working on other tasks.

Excursion: Programming Arduinos

The name Arduino not only refers to the microcontroller boards, but also to a development environment. It features an IDE written in Java, which closely resembles that of the Wiring (thus staying true to its heritage). This environment provides some basic code editing features that are similar to those of many enhanced text editors; it provides syntax highlighting, brace matching and automatic indention. It enables a developer to compile code and upload it onto an Arduino board with the click of a single button; arguably, this is its most compelling feature, especially for novice users. The same process would otherwise be achieved by utilising an AVR tool chain. According to its homepage, the "*Arduino programming language is an implementation of Wiring, a similar physical computing platform, which is based on the Processing multimedia programming environment*" [10]. For running on Atmel processors it links against the library `avr-libc`, which provides a subset of the standard `C` programming language functions. An Arduino program (referred to as *sketch*) requires two functions to be defined: `setup()` and `loop()`. As the name implies, the former is executed once after powering on or resetting the device, while the latter is continuously executed until the Arduino loses power (or runs out of memory, or the code crashes, or the world ends). Since this would not constitute a valid `C` program on its own, the whole sketch is transformed before being compiled and uploaded. Like with other programming languages, there are many extensions to the core language functions, which are called *libraries*. These are written in `C` or `C++` and may be used in the Arduino code via a simple include after being extracted to a predefined folder within the Arduino IDE's workspace (called *sketchbook*).

does provide, however, are *interrupt timers*. Each Arduino board has a set of timers (three is the most common number) that can be set to interrupt the normal program flow once they reach a certain number of ticks. They then execute a user-defined function and once that is done let the `loop()` pick up where it was interrupted. I set one interrupt timer to check the button states and found that this worked quite well, with acceptable response times to button presses.

This is it for the back-end part of the device's software. Now for the user interface. The main defining factor for how the user interface turned out in the end was the small screen. It did not allow for many elements to be visible at the same time, lest they become too small and badly readable; this necessitated that I reduce the information on the screen to a minimum, containing only the most important bits. This is not necessarily a bad thing, as it led to a lot of thinking about what was really necessary for the user to see.

As you have certainly noticed by now, besides the power switch there are three possibilities for the user to interact with the device, all located below the screen: The device's red centre button is always used for activating the currently selected item; this includes entering menus, acknowledging alarms or changing options. The buttons to the left and right move the current selection up and down, respectively. The buttons always retain the same functionality in all of the menus and various screens, which I deemed important for the sake of consistency.

(a) All present

(b) Missing one "Susi Huber"

Figure 3.9: Main screen in both possible states

The first thing the users sees after powering on the device is the main screen, which I usually refer to as standby screen. It contains only three elements: at the very top there is an indicator of the current range setting ("Reichweite mittel", German for medium range, in Figure 3.9). Beneath that is an area indicating the group's status. It may either be green (Figure 3.9a), which means that all is well, or red (Figure 3.9b) if someone is missing, in which case it also displays the names of the members in question; multiple absentee's names will be displayed beneath each other, one per line. Lastly there is a label for entering the main menu at the bottom.

The main menu, as can be seen in Figure 3.10a, consists of six options altogether. Each one leads to another screen, but beyond that there are no further levels; the whole interface is never more than two levels deep. For the sake of translation: the available options in the main menu are, from top to bottom: changing the maximum range, viewing a list of currently registered members, sending a notification to all members, changing general device settings, registering new members and finally going back to the standby screen. I will give a short description of each of these screens and their functions.

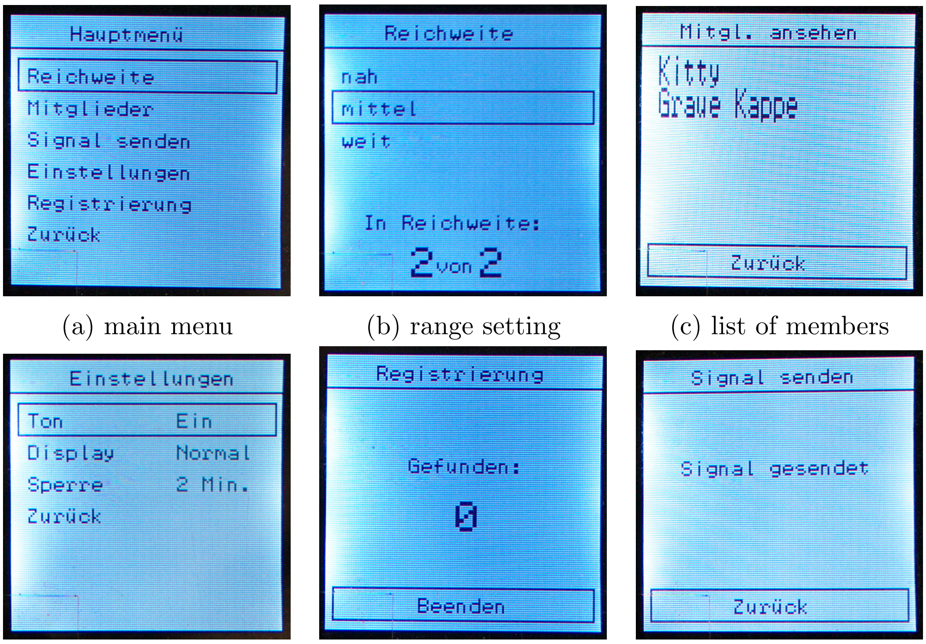

(a) main menu (b) range setting (c) list of members

(d) options (e) registration (f) sending a signal

Figure 3.10: Menus of the main device

Reichweite ("range") leads to the screen shown on Figure 3.10b. In the upper half of the screen, there are three options to chose form, which translate to "near", "medium" and "far", in that order. Beneath that is an information area that shows how many of the currently registered members would be in range if the selected settings was chosen ("In Reichweite" meaning "in range"). That is actually the only mock-up element remaining in the prototype. It was visible only during the first usability test to see whether or not participants would react to it; the number of in-range devices on the last line was hardcoded to the selections and not dynamic.

Mitglieder ("members") shows a list displaying all currently registered members of the group (Figure 3.10c). It also indicates the state of the members; those out of range (for which an alarm would also be sounded) are displayed on a red background. There is nothing to do for the user in this menu, it is purely informative.

Signal senden ("send signal") brings the user to a screen (Figure 3.10f) displaying at first a message indicating that the signal is being sent, which changes to an indication that this has been done after no longer than two seconds (this is the maximum time it takes for the back-end to transmit the requests). There is nothing to do for the user except go back to the main menu.

Einstellungen ("settings") shows the device's general settings. As you can see on Figure 3.10d, there are only three of them: the first one ("Ton", translating to "sound") lets the user choose between enabling or disabling the audible alarm (as there are currently no other sounds implemented). The second option lets the user rotate the display 90 degrees clockwise, thus placing the buttons to the right of the screen. This was meant to provide an alternate way of handling the device, possibly better facilitating operation with one hand only. However I never examined the impacts of this during either usability or field testing. The last option lets the user choose how long the device may remain idle (meaning: without user input) before it is locked.

Registrierung ("Registration") enables a setup period for the group. In this time, all devices that are part of the ZigBee network and transmit a correct request for admission are automatically registered as members. The user is shown a screen that indicates how many additions there have been to the network, which is updated whenever a registration occurs. This setup period is active for as long as the user remains on this screen (which is shown in Figure 3.10e). More information about the setup of the network and various request types may be found in Section 3.3.

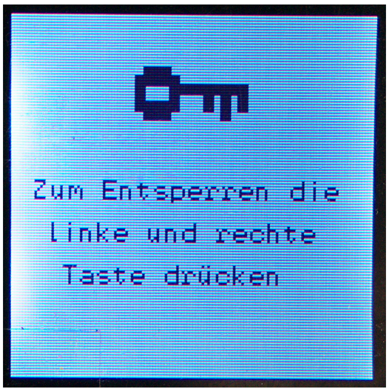

Figure 3.11: (un)lock screen

I tried to keep the layout and usage of the various screens as consistent as possible. The menus are all structured the same: the current selection is always surrounded by a rectangle, even if there is only one possible selection (e.g. on the standby screen, when acknowledging an alarm or when registering new members). The menu that the user is in is always displayed at the top of the screen as a heading, with the exception of the standby screen (where it is replaced by the active range setting) and the unlock screen (where there is none), as I did not consider either to be a menu. If there are multiple options, the last one is always “Zurück” (back), which leads back to the previous screen.

Lastly there is the (un)lock screen shown in Figure 3.11. As I have said earlier, if a certain time passes without any input from the user, the device will turn off the display. Once the user presses a button again, the unlock screen will be shown. It holds the symbol of a key and a short message for the user which says that he or she should now press both the outer left and right button to unlock the screen again. If this is done, the user is dropped into the standby screen. The screen lock is automatically disabled (and the display switched back on) if an alarm occurs.

You have now seen all of the different states, screens and menus which the device has to offer. Apart from that, there are no other possibilities for the user to interact with the main device.

Improvements After Usability Testing

After I had conducted the usability tests with the device as introduced above, some flaws became obvious and I made the according changes to the prototype. These shall be discussed in this section so that they may be compared to the original version of the device. Please refer to Section 4.1 for a detailed description of the usability tests. The changes listed here were done prior to conducting any field tests with the device.

(a) new and old buttons in comparison (b) exterior

Figure 3.12: New buttons for the main device

The first and most glaring flaw in the interactions with the prototype were the buttons; not the fact that there *were* buttons, but rather *which* type of buttons was used. I admit that neither their look nor their feel did much for the device. They protruded visibly from the surface, making it harder to store the device in pockets or bags. They also needed to be firmly pressed all the way down to register an interaction; considering their lift-way, this made them cumbersome to use. Therefore I decided to completely replace them. The new buttons were more flat, meaning that they would not protrude so much from the device, making it less likely to catch on something while being carried around. This also means that the user would not have to press the buttons down so far, as their lift was lower. Lastly they make an audible clicking sound and give tactile feedback once the threshold for registering the button press is reached. Overall they felt much more solid and comfortable to use. Figure 3.12a shows side-by-side comparison of both button types.

The buttons were glued onto the casing's inside and re-wired, visible on Figure 3.14. Since the participants of the usability tests had shown good understanding of the button's functionalities, I painted the centre button red to mimic the colour coding as it was. Figure 3.12b shows the device with the new buttons in place.

(a) main menu

(b) member management

Figure 3.13: Menu after the introduction of a new members management screen. The option "Debug" in 3.13a was not visible during field testing

On the software side of things, I only made some subtle changes. Most notably is that I introduced a dedicated members option on the top level of the main menu. The options for registering new members and viewing a list of current members were moved in there, together with a new option to remove single members from the group. This was the first and only structure of the user interface that is more than two layers deep (the only "menu within a menu", if you will). The screens for member registration and listing remained the same, however. The new menu and its top level entry may be seen in Figure 3.13a and 3.13b, respectively. I also changed the wording of the lock screen. The word "gleichzeitig" (meaning simultaneous) was added to indicate that the buttons should be pushed at the same time. I also changed the label of the buttons in both the settings and registration screen to say "Save" instead of "back" or "complete".

These were all the changes I considered necessary after the usability tests. Additionally I did some re-wiring of the device's internal components; this was purely for cosmetic reasons, however, and did not influence the functionality. You can see the internals on Figure 3.14 in the state they remained for the rest of the project (which is to say, two field tests).

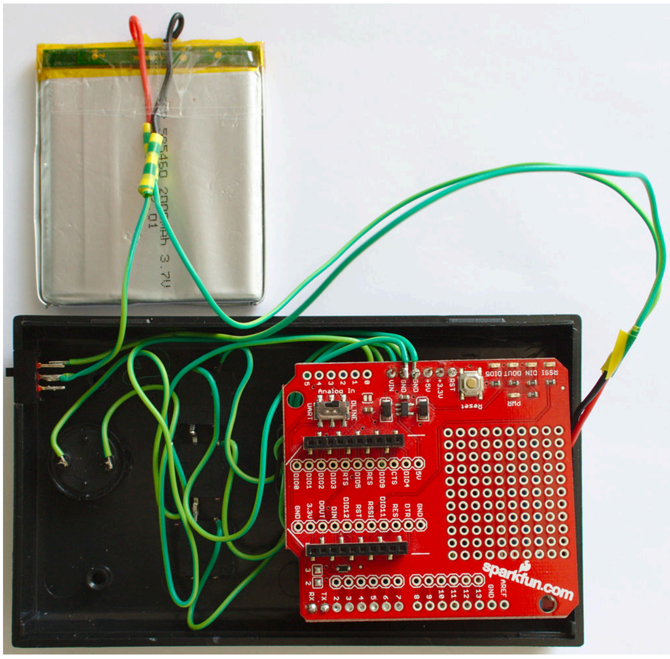

Figure 3.14: Main device, revised internal wiring

3.2 The Child Devices

After having implemented, evaluated and then improved the main device prototype, I turned towards manufacturing the children's devices in preparation of the first field test. As I have already mentioned briefly on page 44, the Arduino prototyping platform also features a brand of devices intended for use with textiles: the *LilyPad*.

The LilyPad is fully compatible with the Arduino project, both in hardware and software. However, its design was adapted such that it could be easily fastened to clothing and connected to other components through conductive threading. The board is in the shape of a circle and kept quite flat.

It has multiple holes placed around the edge in even spacing. These may be used to sew it onto textiles, but more importantly they also serve as contact points for the LilyPads input and output pins. It is very easy to fasten a conductive thread through these, which may then be connected to other devices.

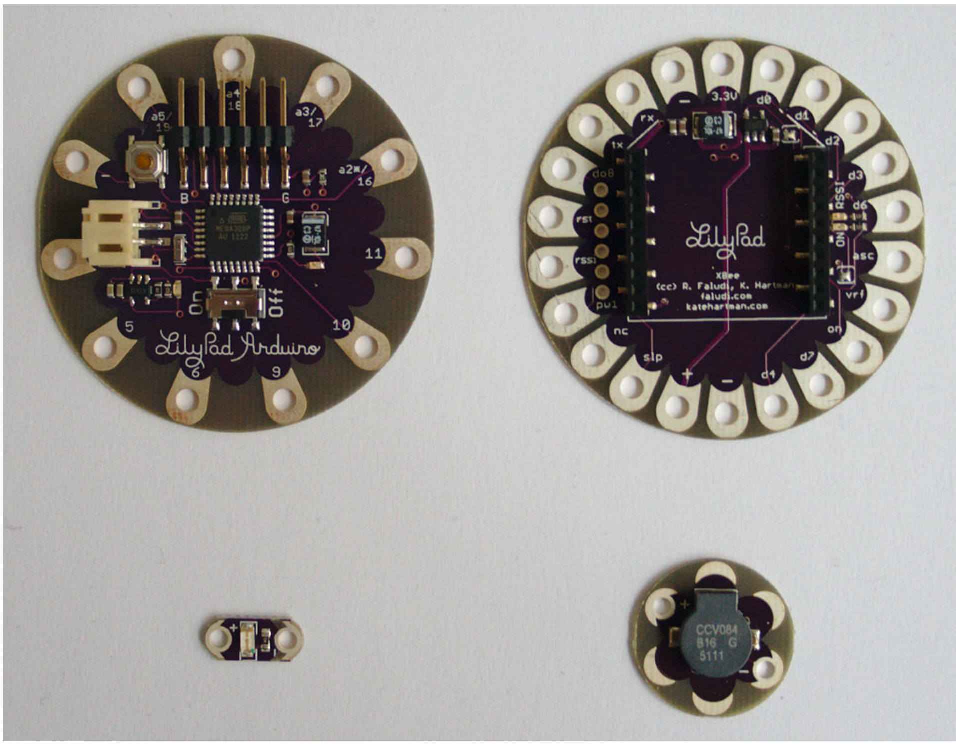

Figure 3.15: Electrical components used for the child devices, clockwise from top left: Lilypad Simple Board, LilyPad XBee Shield, a buzzer and one LED light

Like with regular Arduinos, there exist several flavours of main boards for the LilyPad too. I have opted for the *LilyPad Simple*, which is shown in Figure 3.15 (top left). It is based on the same microprocessor as the Arduino Pro used in the main device, the ATmega328. Unlike its sibling, however, the LilyPad runs on a lower input voltage (2.7V - 5.5V) and consequently provides a lower output voltage also. This furthermore means that the board runs at a lower clock speed, namely 8MHz, which is half of what the Pro provides. These features are common to all LilyPad variants currently on the market, by the way. Apart from that it features 32KB of flash memory and 2KB of SRAM.

I chose this particular LilyPad board because it features a JST connector ex-factory, which means I can use the same Polymer Lithium Ion battery packs that are commonly used with Arduino projects. It also provides a charging circuit for said batteries, meaning that no additional external devices will be needed for recharging. Lastly it also comes with an integrated, onboard on/off switch. All of these are features that are not found on the original LilyPad. The only downside is that the Simple board has fewer pins available; 9 instead of 20 to be precise. This is not a problem at all for my designs, which only use two of the pins (three if a buzzer or LED light is added for notifications). Arduino's LilyPad brand also features an XBee shield with the same form factor as the main board. Figure 3.15 shows the shield next to the main board on the top right.

Apart from these two main parts there was also one buzzer (Figure 3.15 bottom right) and several LED lights in red colour (one is shown in Figure 3.15, bottom left). Both are meant for notifications on the child devices, as they were wished for on multiple occasions during the requirement analysis interviews. Each prototype for the children would also include a smaller version of the battery pack used on the main device, providing 1000mAh (half of what the bigger variant achieves).

(a) aimed at boys (b) intended for girls

Figure 3.16: Headwear to hold the LilyPad components

At this point in the project I needed to do some clothes shopping. I decided to try various forms of wearables as hosts for the prototypes. Altogether I had enough parts to manufacture five child devices. The preliminary interviews had already given me some ideas about what could work. First I bought two high visibility vests; one cannot really go wrong there, they are basically all the same. Then I purchased two headwear items: first a cap meant primarily for boys (Figure 3.16a) and another one which I hoped would appeal to girls (Figure 3.16b). For the last device I wanted a pendant; this would later be hand-crafted completely, as I did not find anything ready-made that seemed suitable.

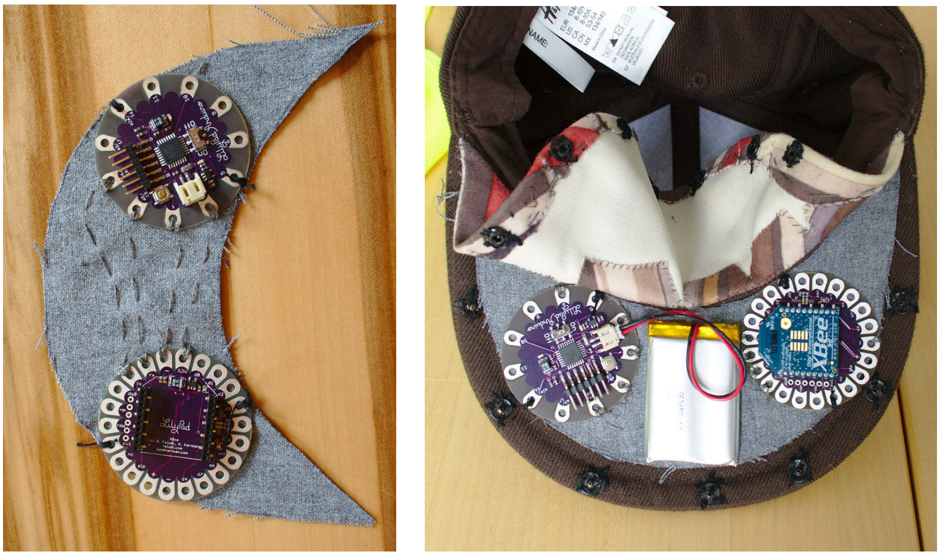

(a) (b)

Figure 3.17: Adaption of the brown cap to house the LilyPad components

With the items acquired it was now time to integrate the LilyPads. As no two devices would be the same in the end, this turned out to be quite some work. First let us take a look at the headwear items. The brown cap actually seemed to provide enough space on the visor to support all of the LilyPad components. They were sewn onto a patch of fabric cut out in the shape of the visor (Figure 3.17a). To hide them from sight and provide some minimal protection, an additional layer of cloth was sewn to the visor and fastened with some snap buttons at the front; this was to allow easy access to the electronic components later in the project. Figure 3.17b shows the final version.

The other headwear did not provide such ample space on the visor. Therefore I opted to integrate the electronics into the inside of cap. To this end the LilyPads were again sewn onto a separate patch of fabric (Figure 3.18a) that was then sewn onto the cap on one end and fastened with Velcro on the rest of its length. The result can be seen in Figure 3.18b.

So much for the headwear items. Next came the high visibility vests. Their setup was identical, except that one of them featured a buzzer and the other did not; for illustration purposes in this section, all photographs will be of the former item. The LilyPads were sewn directly onto the vests in the area of the upper chest, near the shoulders (Figure 3.19). A patch of tougher fabric was sewn onto the backside of that area to provide some stability and prevent short-circuiting through skin contact or perspiration. The battery was fastened on the other shoulder to provide some form of counterweight and symmetry. The wire to connecting it to the LilyPad was

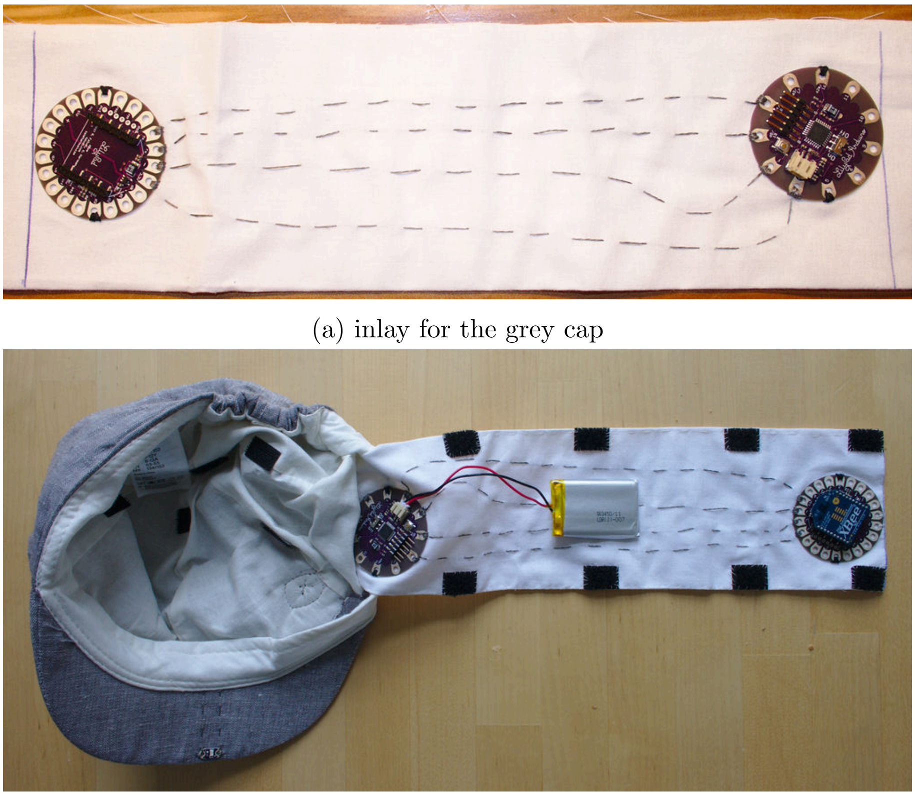

(a) inlay for the grey cap

(b) finalized internals of the grey cap

Figure 3.18: Adaption of the grey cap to house the LilyPad components

worked into the collar of the vest, as can be seen in Figure 3.20. As a last step, fabric (obtained by sacrificing a differently coloured high visibility vest) was sewn over the electronics to provide some protection (Figure 3.21a). It was fastened on one end with Velcro only to facilitate easier access to the LilyPad for later adjustments (see Figure 3.21b).

The last of the devices was the pendant. This was the easiest to manufacture: the electronics were sewn onto a patch of cloth that was itself glued onto a congruent piece of cardboard (to provide some stiffness), which is shown in Figure 3.22a. A fabric casing was fashioned to fit this (Figure 3.22b). As with the other items, it was designed such that subsequent access to the electronics was possible. This was achieved by closing it on one side with snap buttons only, such that the insides holding the LilyPad could be removed completely with ease.

As the preliminary interviews had revealed, there was some demand for a way to notify the children of important events, like exiting a public transport or meeting at a gathering place. To this end, three of the five devices received some form of notification feature. On the caps, this was

Figure 3.19: LilyPad components on a high visibility vest

Figure 3.20: High visibility vest with added power circuitry

done in the form of a LED light, integrated into the very front of the visor; one of them can be seen on Figure 3.18b. I opted to integrate a buzzer on the high visibility vest for an auditive alarm signal; you can see it on Figure 3.19; it is the smallest device near the bottom. The LED lights were visible if the caps were being worn even halfway as intended, since they proved surprisingly bright even in daylight. The buzzer on the vest was also well

(a) fabric to cover the LilyPad (b) a peek at the electronics beneath

Figure 3.21: High visibility vest, final version

audible. Yet whether or not the children actually noticed them remained to be seen during field testing.

That is it for the physical part of the LilyPad devices. The software side is actually quite simple. In the `setup()` method, the XBee module is initialised together with a software serial interface[4]. The XBee module is queried to obtain the node identifier, which the LilyPad will use as name in subsequent communications with the main device (see Section 3.3 for more on that). As a last action in `setup()`, the buzzer is sounded (if available) or the LED made to blink (if available).

The `loop()` first checks for incoming transmissions with a timeout of 500 milliseconds. If a message is received, the request type is retrieved and the according action taken. Possible incoming messages are: a status update from the main device (sent periodically), a notification indicating that the device should now use either LED or buzzer (whichever is available) to notify its wearer and lastly several possible responses indicating either success or failure after the LilyPad tried to register or report in with the main device. Once that is done, the device checks whether it is time to report in with the main device (configured to happen every two seconds during the field tests) *if* it has previously made a successful registration attempt. Otherwise `loop()` tries to do another registration request with the main device if a certain time has passed since the last one. These are all of the primary responsibilities

[4]the LilyPad does not feature a dedicated serial interface and UART like most Arduino boards; therefore a special library, SoftwareSerial[15], is used to emulate it on two input/output pins.

(a) electronic innards (b) pendant casing

Figure 3.22: Pendant child device

of the `loop()`. In early versions of development I had also made it sleep for several seconds if no network association was indicated by the XBee module; this was disabled after I discovered that an XBee router (which all of the child devices were configured as – see Section 3.3) always returns a positive association indicator, rendering the functionality useless. As it stands, no sleep or power saving cycles were implemented.

3.3 Networking It All Together

As I have established and explained in Section 2.2, I believe that ZigBee is the most fitting wireless technology to use for my project. One of the reasons for choosing Arduino as prototyping platform was that it provides an implementation of the ZigBee standard, called XBee. In this section I will describe my wireless network setup, which is based on said protocol.

The hardware side of XBee, as you likely have noticed if you have read the previous sections, is split into a wireless module and a shield that serves as an integration environment with the Arduino. The type of wireless modules which I have used on all devices is shown in Figure 3.23 - it is a "Series 2" version of the XBee modules sold by Digi. "Series 2" in this case indicates that it is capable of real mesh networking and also features a true implementation of the ZigBee standard. Prior versions did not have these features, relying on a proprietary protocol to achieve a semblance of mesh networks. The wireless shields for the main and child devices can be seen on page 3.3 and 3.15, respectively.

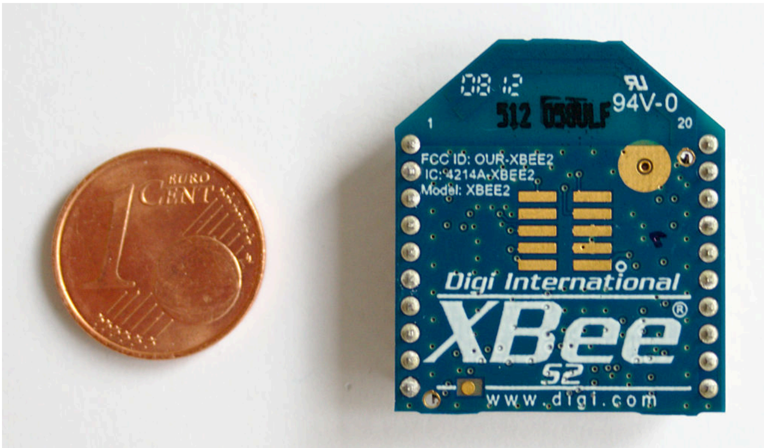

Figure 3.23: XBee module as was used on all devices for this project

The configuration of the XBee network is done and stored completely on the wireless modules themselves; they have a persistent memory of all relevant parameters. Programming the modules may be done through an application provided by the manufacturer, called *X-CTU*. For connecting the modules to a computer I have used a device called "XBee explorer", depicted on Figure 3.24, which allows for the radio to be plugged into a USB port. The XBee modules could also be programmed directly through an Arduino by writing a small sketch; I found the former method to be more convenient.

There were not that many configurations to be done in order to set up a network between the devices. The "Pan ID" parameter, which indicated the ID of the network, had to be set to the same value on all devices that were to participate; the actual value of the parameter is arbitrary, it is merely required that all radios use the same Pan ID. API mode was also enabled on all of the devices, since it is required for the XBee Arduino library to work. On XBee modules, API mode means that an alternative, frame based transmission method is used. It is more advanced than the standard (often called "transparent") transmissions, offers more possibilities

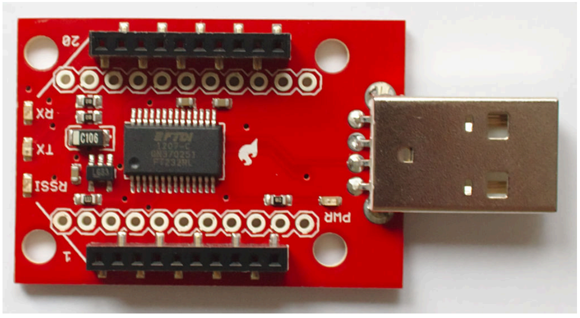

Figure 3.24: XBee explorer dongle used for programming the radio modules

(like viewing the source address, LQI or delivery status of a packet), but also more complex (in the way that each packet is constructed). Lastly the node identifier parameter was set on each module. It holds a forty character string which serves as the device's "name". This was very accommodating to my needs, as it allowed me to set the children's names via the XBee module configuration as opposed to having to do that programmatically in an Arduino sketch (which would have resulted in different sketches on a per-device basis). No other configuration was done to form the ZigBee network. For those of you that are interested in the details: all of the modules were flashed with firmware XB24-ZB, revisions 2x64.

ZigBee's device types and topologies

ZigBee differentiates between three different device types [64]. Each network is managed by exactly one *coordinator*. Furthermore there is an arbitrary (or at least big enough as makes no difference) number of *routers*. The coordinator may extend the network through such routers, thus they form a backbone to the whole infrastructure. Lastly there are the simple *end devices*, with reduced functionality but also features like power saving cycles. The coordinator and routers each have a limited number of possible children that route their communication through them; these may either be routers themselves or child devices. Using those three device types, the ZigBee standard defines several network topologies [28]. The most simple is the *star network*, where each end device is connected only to the coordinator. Then there is the *tree network*, which is similar to the eponymous construct known from graph theory; the routers represent branch nodes and the end devices form the leaves. Last there is the most complex topology, the *mesh network*. In this setup, each device may be connected to any other, so long as the latter has routing capabilities (i.e. is either the coordinator or a router).

The ZigBee coordinator module was integrated with the main device prototype. All child devices received modules configured as routers during the field tests. The main reason for this was that I wanted to test the mesh networking capabilities. There is no inherent order in a group of school children, their position is constantly shifting. Therefore it would not be practical to hand out a mix of routers and end devices: it may well be that all of the routers end up on one end of the group, the coordinator on the other and the end devices somewhere in between. There is no possibility to be certain that the different device types distribute evenly, hence the solution I chose was to make them all equal.

Another reason for configuring all of the radios as routers was that I was not able to produce satisfactory results with end devices. First of all, as required by the ZigBee standard, the routers have no sleep period. They are always powered on and always immediately send and receive. End devices, on the other hand, have a sleep cycle during which their parent (either a router or the coordinator) keeps their messages for them and retries sending them several times in the hope that the receiver has woken yet. There were performance problems with all modules (apart from the coordinator of course) being configured as end devices. The rate of dropped messages was noticeably higher with this configuration, resulting in the main device reporting many false positives for missing devices. This observation is backed by [64], where it is revealed that, at least in mesh networks, end devices perform worse than routers in both sending and receiving, which is especially true for the mobile context. Thus I opted to go with what I knew would work, an all-router configuration. More on this topic can be found on page 70 and following.

The actual functionality of the network, to assert that all of the members are within range, is pretty straightforward and, I have to admit to my dismay, not very sophisticated. Let us start with the registration. After powering on the main device, it knows nothing of any group members, as there is no persistent storage integrated in this prototype. The ZigBee network itself is configured to allow joining at all times, thus any device configured with the correct PAN ID setting would be accepted.

However, the main device maintains its own list of admitted group members that it watches out for, which constitutes an additional administrative layer above all of the ZigBee stack. While the registration period is active, the child devices may send a request for registration, which also contains their human readable names (as taken from the Node Identifier parameter of the XBee module). Any node that is part of the ZigBee network and sends a correct registration request to the main device during the registration time window will be added to a list on the main device. To be more precise, there are three of them: one list to match names to an address, one to track the last time an update was received per address and one experimental that keeps track of the last three RSS readings per address.

Once a network node has registered with the main device, it is expected to send regular update requests to the latter. Since I have mentioned certain types of requests several times now, here is a short explanation of what I mean by that: each message sent between the devices of the network follows a template, with the third index always carrying the indicator for the message type (the first and second index being a message delimiter and padding, respectively). Both the main and child device always look at the third index of a data packets payload first to determine the message type and act accordingly. The only request type that has any additional data, aside from the message type, is a registration request sent by a child device; it also carries the latter's name.

If the main device does not receive any update requests from a child for a certain period of time (which was set to six seconds during the field tests), it will sound an alarm. There is nothing more to the presence detection capabilities of the network. As I have said, this is not very sophisticated, but it works; unlike some more complex networking setups I have tried but failed to implement. But more one these in the next section.

Problem Scenarios One and Three, or: This is Harder Than I Thought

I had several plans how to tackle the problem scenarios one and three. To remind you, these were the boarding and exiting of public transport vehicles and the transit through highly frequented areas, as discerned in on page 12 and following. Sadly, none of my approaches worked out in time to be implemented for field testing.

As I have stated in Section 2.3, it is not possible for me to discern the position of any node within the ZigBee network that the prototype devices form. Any form of localisation needs a certain number of beacons, devices with known positions, to be performed. I do not have those at my disposal. The exception from this rule is the approach presented in [105], which establishes relative positioning without the need for beacons; yet it utilises time of arrival distance measurements, which I cannot provide with the Arduino hardware (explanation follows). With that established, I wanted to provide what I knew was, in theory, doable with the hardware I had: range estimation. Range estimation would be a huge benefit to the prototypes, as it would facilitate on-the-fly definition of how far the children are allowed to venture from the teacher. Checking for absentees when, for example, using the public transport would require a much more restrictive range to quickly notice stragglers than during a stay on a playground. So to tackle more than one of the three problem scenarios I had discerned, it was essential that I be able to gauge how far one node in the network (a child, in this case) was removed from the network “sink” (the teacher).

My first investigation was into the realm of time based distance estimates. Time of arrival (ToA) was out of the question since the beginning, as it requires exact and synchronised time keeping on all devices. There is also the technique of measuring the round trip time (RTT) of a transmission, which eliminates this need. However, there is a fundamental problem with using time based estimates on my prototype: the hardware does not provide for it. Radio waves move at the speed of light, which is 299 709 km/s in air. Put differently, a radio signal takes about 3.3 nanoseconds to travel one metre. The Arduinos I use (and pretty much any other Arduino board as of the time of this writing) run at 16MHz and 8MHz. That means one computation cycle every 62.5 nanoseconds. If you do the math, you will see that such processing power does not provide the resolution needed for time of flight measurements by a long stretch. Additionally, the XBee modules themselves take some microseconds at the least to reply to any request with an acknowledgement; after doing some quick tests it seems that the unknown reply time of the transceiver varies in the range of several microseconds, which makes it completely useless for this form of applications. To make it short, time of flight based distance measurements are impossible with the hardware available on this project; it’s about forty Euros worth of generic equipment, after all, not any specialised high precision instrument.

However, the XBee modules also provide RSSI readings which I could use for distance estimations. The IEEE 802.15.4 standard (on which the ZigBee protocol stack is based) states in its 2011 revision that “*LQI measurements shall be performed for each received packet*” [1]. It defines LQI as “*a characterization of the strength and/or quality of a received packet*”, with possible values ranging from 0x00 (0) to 0xff (255). It further demands that all values between this minimum and maximum be uniformly distributed, with at least eight unique values being provided. According to the manual of the ZigBee modules [2], such readings are provided and may be retrieved through a special command. The observed values of these readings lie between 0x1A and 0x5C; this is not really in accordance with the 802.15.4 standard, but in lieu of alternatives it would have to do. The biggest problem was not the range of the values, but rather that they did not provide “accurate” (according to the manual) values for multihop links. It is not explained any further what “accurate” means in this case; I assume an indicator is only ever provided for the last hop of any transmission.

My first approach to utilize the signal strength indicator provided by the XBee modules was to add an additional layer of management to the network. This was necessary because, as mentioned above, the RSS reading only referred to the last hop of a transmission. Since I had configured the network such that it could, in theory, have multi-hop transmissions, this was the first problem to tackle. My plan was to let any message “bubble up” through the network topology until it reached the coordinator, appending each hop’s signal strength reading to the data payload as it went along. However, this proved to be more of a challenge than I had expected. The XBee modules support a command that returns the address of their parent, which seemed to be ideal for what I had in mind. After some experimentation it turns out that this only returns valid results for end devices, not routers (of which the whole network spare the coordinator was consisting at that point).

So obviously that was not the right angle. Since the date of the second field test was quite near, I tried a little low effort trick to get at least some form of RSSI measurement: I configured all of the child devices to work as ZigBee end devices, effectively forming a star topology centred on the coordinator. This did indeed enable a viable reading of signal strength values at the main device. I implemented a very simple moving average system: the last three readings from any device were kept in memory and at the periodic checks for absentees their mean was calculated and compared to a threshold; the latter was defined by the current range setting. Initially, this approach yielded some promising results. In the end, the system proved much more unstable during testing than before. The packet loss ratio was much higher, resulting in the main device reporting missing group members when they were in fact only a few feet away. Since the algorithm for tracking and storing signal strength readings was not computationally expansive and sufficient memory was still available, I attribute this to the end device configuration. Indeed, work done in [64] confirms that end devices perform worse than routers in mesh networks.

Another problem I encountered was that the RSSI readings always referred to the last received packet. This did not necessarily mean the last update from a child device; it could also be an acknowledgement frame or network management overhead packets that were not explicitly sent by my code on the child devices. As I have stated the main device listens for incoming packets in between performing display updates and record keeping; I took the signal strength measurements immediately after receiving a packet, yet sometimes I would find inexplicable values for devices that were way larger or smaller than their last few readings.

Ultimately, these tweaks to incorporate RSSI readings into my system only added an air of instability that I was not happy with. The results during "laboratory" testing[5] were less satisfactory than before and it also meant sacrificing mesh networking capabilities. As I had not managed to improve on these areas shortly before the date of the second field test, I opted to revert back to the proven setup. Sadly this also means that no range estimation was tested in the field with this prototype and is indeed not present in the final version of the software at all. Much as it irks me to admit, I have to say that I clearly failed in the goal of implementing a

[5]read as "at home on my work desk"; but that just sounds so much less academic than "laboratory"

distance estimation system in the course of this work.

About Security

As you may have noticed while reading this section, the network setup used by this prototype is pretty far removed from any definition of “security”. I am aware of this, and the fact that the situation could be greatly improved even with what the prototype hardware provides. I am at this point quite loath to strain the phrase “beyond the scope of this work” any further; yet the security details are exactly that, for the reason of time constraints. None of the security features that are part of the ZigBee standard were utilised in the prototypes. Thus there are several vectors of attack that a malevolent

Excursion: Security Measures Provided by ZigBee

ZigBee offers a set of possible security measures, affecting the MAC, network and application layer; they include data encryption, integrity and authenticity checks as well as data freshness assurance and are based on an 128-bit AES algorithm [63, 36]. ZigBee knows three types of keys to provide these functionalities [95]: the *network key* is known by all devices on the network and required for communications between them, the *link key* protects point-to-point communications between two devices and the *master key* serves as a shared secret to derive the two former keys from. *Integrity* is assured through a message integrity code provided with each data frame to detect possible changes to the packet. *Authenticity* checks are done via a modified version of AES-CCM (Counter with Cipher Block Chaining Message Authentication Code, or CBC-MAC), which supports either encryption-only, integrity-only or both. *Authentication* may be provided in three different ways: first an Access Control List of authorised nodes may be maintained on each device, necessitating authentication via the MAC address. Secondly, a standard security level trust centre may be utilised to grant access to new nodes by issuing them a network key (in plain text), which is then required for all further communication; alternatively, new devices may have a network key already installed. Lastly there is the possibility of a high security trust centre, in which case joining nodes must already be aware of the network's master key and, after doing a four-way handshake with the coordinator to assure the validity of the key, may derive the network key from it. ZigBee assures the *freshness* of data by maintaining a frame counter on the communicating devices, which is updated every time a new frame is received; data packets containing an outdated counter value are discarded.

actor could use to bring down the whole system or distort the actual state of its members. The only thing preventing an attacker from being part of the ZigBee network is not knowing the correct PAN ID. Since this configuration parameter is neither hidden (nor meant to be) and its discovery is indeed an integral part of the ZigBee protocol, it may be obtained through sniffing quite easily [36]. Therefore it would not be difficult at all for any ZigBee device to become part of the prototype network. Once that is done the possibilities for disrupting regular operations are myriad. Lets indulge in naming a few:

Denial of service an attacker could simply spam the main device with data packets, rendering the whole system useless. The regular updates from the child devices would not get through at a frequency that makes any difference, because the coordinator is busy dealing with a flood of other packets.

Impersonation it would be possible to impersonate any child device by replaying its status update messages to the coordinator. The latter could be kept from realising someone was missing indefinitely, so long as the "fake" child is still within range. More importantly, an attacker inside the network could also impersonate the coordinator and send networking commands to the child devices.

Privacy breach during the registration process, the Node ID parameter of the ZigBee module used with a child device is sent to the main device as clear text; it is meant to contain the child's name, or at least some form of label by which a teacher may unambiguously recognise it. Since there is no encryption employed in the network, an attacker may obtain this information without much effort.

With that all being said, I have to state that, in my opinion, security should not be a primary issue in any further work on this subject either. My reasoning is this: in none of the application scenarios I have described in Section 1.1 would hacking this network be worth the hassle. I see no benefit to be gained for a malevolent attacker to bring down the system. There is little private data to be sniffed, no malware to be installed on the devices, no location tracking to compromise the privacy of any users. The only sensitive information that is ever sent through the network is the name of the children, once, during registration; the registration process will most likely not be done in public places, but rather during setup before venturing out.

Admittedly, as humankind has proven time and time again since we climbed down from the trees, there will be those that aim to disrupt something purely because they take joy in it. These would indeed find easy prey in a system such as my prototype. And yet, even if I had utilised all of the security features provided by ZigBee, there are much less subtle ways to cripple this network than to hack into it. As indicated on page 34, ZigBee devices use low power radio transmissions and are susceptible to interference; a powerful radio transmitter used with the purpose of disrupting a ZigBee network may be enough to bring down regular operations. Brute-forcing a degradation in signal quality to a certain point would most likely render the system defunct, with no hacking required.

Chapter 4

Testing and Evaluation

In this chapter I will describe the trials of my prototype to evaluate its usability and usefulness. I will explain how these were conducted, what I could observe during their course, what my conclusions are concerning these observations and what consequences should be drawn from them. To this end I have conducted three different tests: Firstly there were usability tests under laboratory conditions, which aimed to assess the main device's user interface and judge its readiness for further examination; Section 4.1 holds a description of these tests. The aforementioned "further examination" would be the field tests, of which I conducted two. Their aim was to evaluate the devices usefulness in the field, to reveal usability flaws that may have not come to light under laboratory conditions and to observe the pupil's reactions and behaviour towards the child devices. A report of these tests is given in Sections 4.2 and 4.3.

4.1 Usability Testing Under Laboratory Conditions

After the preliminary interviews as well as my technology evaluation were done, I started implementing a prototype for the main device (more information on this endeavour may be found in chapter 3). It already featured all of the functionality and interfaces I had planned to implement. In this phase of the evaluation I wanted to see how well the intended audience, the teachers, would handle this device. [73] served me as the main guideline for these usability tests.

At this point in the project, the prototype was very much functional. All physical buttons and switches were operational, as was the sound module and of course the display. Software-wise the interface was far advanced as well with all the menus being fully implemented and, apart from the dialogue for changing the range, functional. The device was in a state that could already be used in the field if no issues were discovered in usability testing (which I deemed a highly unlikely event).

Nielsen distinguishes two kinds of evaluation, the *formative* as well as the *summative* one [73, p. 170]. My tests were to be conducted as the former type. Their purpose was to find and specify necessary improvements to the device as part of an iterative design process (which had also involved preliminary interviews and would in the next iteration include field testing). In contrast, summative evaluation would mean to assess the interface in comparison to at least one alternative in order to find the more viable one. As there is currently nothing to compare the interface to, this is not applicable.

Test Plan

Before conducting my usability tests, I defined a test plan which specified the purpose and goals of the tests, where they would take place, which tasks they would involve, how I would explain the tasks to the user and so forth. This section contains a short description of the test plan as defined prior to the actual tests.

The goal was to gauge (for the first time) how well the users handle the device and its interface. I wanted to know whether the handling is intuitive, the information provided to the user is understandable and how well they navigate the interface. However the testing would also extend to the physical aspects of the interface. It was important to know if the user approved of the device's haptic features: whether it was too heavy, if all the buttons are within the user's reach, if the readability of the display was good, if the noise level of the buzzer was sufficient and so forth. In the end, what I wanted to achieve was knowledge about whether the device was ready to be used in the field or what could be done to make it so.

The test persons would be four teachers out of the same pool as I had already used for the preliminary interviews. Therefore all participants were familiar with the project and had already worked with me once. This also means that all of the testers are part of the target audience and thus very representative of the intended users of the device. I took care to select testers from different age groups. There was one person in their early twenties, one in the late twenties, one in their early forties and the last test person was a bit more than fifty years old. I was hopeful that this would reveal different flaws in the interface through the different approaches to technology that these people would surely have. This was as much a diversification as I could make within the target audience.

I would gather participants through the same channel as with the preliminary interviews: my contact within the school approached them and ask whether they had time to test the prototype I had engineered. As all of the possible candidates were already familiar with the project, no further introduction to the matter was necessary. It was not difficult at all to find participants; during the preliminary interviews, the majority of people had asked on their own initiative to test the device. Similar to the preliminary interviews, the tests would take place in the class room of each teacher's respective class - shortly after the children had left school for the day. The reasons for this were twofold: firstly, it would mean minimal effort on the test person's part. They were at that time already in the room and would just have to stay a bit longer for the interviews, thus minimalising their efforts for partaking in the tests. Secondly, it meant that they would be tested in a familiar environment, their own workplace. My hope was that this would help to reduce anxiety. Figure 4.1 shows an example of such an interview location, which is the teacher's desk in one of the class rooms.

Figure 4.1: A teacher's workplace in the classroom

The test tasks were designed such that they covered all of the device's workflows and interactions; they represented all the use cases which I could conceive for the device during actual usage. I took care to give the tasks a certain "flow". By this I mean that, all together, the tasks formed a meaningful interaction such as could be observed in the field. For example one task would be to register some child devices, then the next would be to find out their names in the respective screen. I did not want to include trivial tasks such as "press button one to get to menu A" or similar. For myself, I had prepared a sheet with all the test tasks in the order they would be performed. Each task featured a precisely defined starting point as well as a state in which it would be considered as done. The tasks were formulated so that the end result of one would always be the starting point for the next. This way, the tester remained in control of the device for the whole duration of the test without me having to fidget with it to make the transition from one task to the next.

The task descriptions would be provided to the user in writing only; I prepared each on a separate card, which would be handed to the tester one at a time at the right point during the testing. This has several advantages: firstly, all of the users are guaranteed to receive the same task descriptions and will not be influenced by my oral description, which could vary from case to case if I was careless; that alone could affect the outcome. Secondly, the participants would not have to remember what their task was, as they could always look it up on the written card in front of them. Lastly, by providing the cards one at a time, the user should never get confused as to what his current task was, and neither could he read ahead to anticipate the next one. Figure 4.2 shows a sample of these cards.

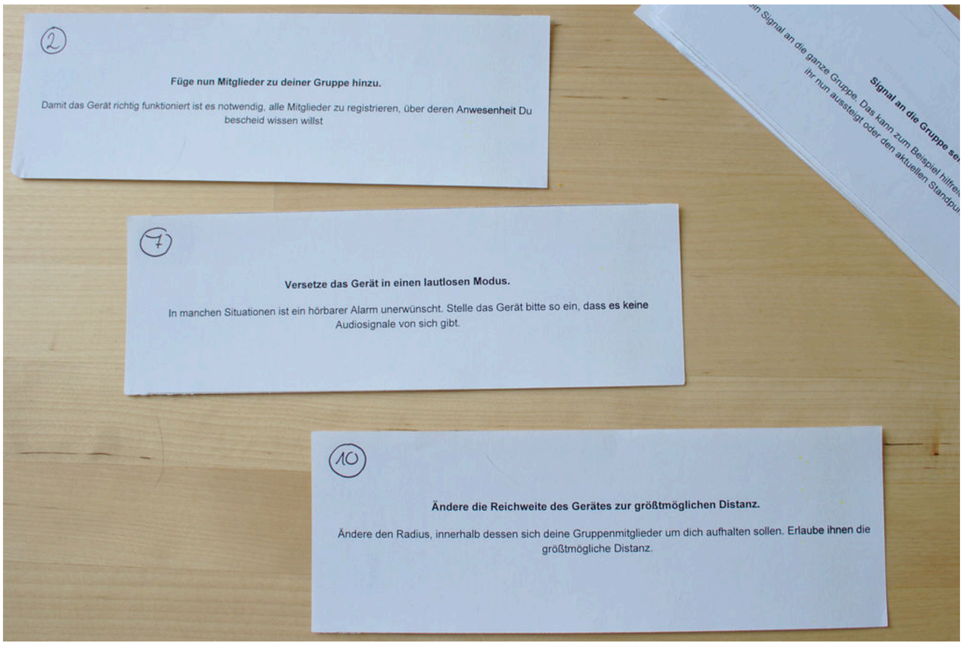

Figure 4.2: Paper cards with task descriptions

As there are no other people involved in the project, I would of course be doing the experiments myself. At the start of each test, the system would be powered off. Since there is no persistent storage implemented in the prototype, this implies default settings and no registered devices. Apart from the main device, two Arduino Fio boards would be present and powered on to serve as “Children” for the duration of the test. No manuals would be provided to aid the user. The system is not overly complex and I hoped to design an interface that can make do without a manual or online help. I planned to intervene only if the user was obviously stuck with a task for more than half a minute. In this case I would provide hints as to what the next step might be. Should this prove insufficient, the task would be considered as failed. I was rather lenient with defining what constitutes a success. I figured that if someone could handle the device without manual or help from my side, this would be a good result. On the note of performance measurement, I would refrain from actually taking times for single tasks. The overall time needed to complete all the tasks will be taken into account, but not as the one defining factor of performance. Rather I would observe and rate how intuitive the interface comes to each person, how much good or bad things they had to say about the device, whether or not they became obviously frustrated with the interactions as well as their general understanding of the task they had just accomplished. I explicitly did not want to set time limits for the users, as I felt that the later field tests would provide enough time-constrained situations.

After the test itself was done, I would ask the participant some additional questions about the device and its usage. These questions involved the brightness and size of the display, the size and weight of the device itself and the sound level of the buzzer. Further I would ask the users their opinion of the physical buttons as opposed to a touchscreen - whether or not they would like a touch based solution better than physical buttons. For gathering data, I would rely on audio recordings and taking notes. I refrained from using video recordings. The reason for this was that the device was small and did not offer much in the way of physical interaction apart from operating one switch and three buttons. The screen was too small to be properly captured with a video camera. Thus I assumed that there was little to be gained from taking a video recording, apart from making the testers nervous.

Testing

Prior to the actual usability tests I conducted one pilot test. The aim of this pilot was to examine the viability of the testing procedure itself, instead of examining the device. The subject for this pilot was not of the pool of intended users, but instead a person in their late twenties which was more easily available to me. Running this pilot test revealed some unclear wording in my task descriptions, but did not show any severe flaws in the test plan itself. Thus I did not change the tasks or the procedure.

The actual usability tests were done in the time between April 7^{th} and 9^{th}, 2013. I would arrive at the school about fifteen minutes prior to the scheduled start of the test and talk to my contact at the school. At the appointed time she would show me to the class room of the day's tester. At this time, the class room would already be empty except for the teacher. The environment was suitable enough for usability testing - there were no outside distractions, the room could be closed off by locking the door to the hallway and I presume that the test users felt at ease; it was their day-to-day workplace, after all. Before beginning with the tests I needed a few minutes time to prepare the device and my documents and recordings. Ideally, I should have been in the room beforehand to ready everything so the testing could commence immediately, yet for organisational purposes (i.e. the test user being in the room already or entering it with me) this was not possible. At the start I would welcome the test user and thank them for participating in the testing. I would further explain to them what progress the project has made since we last spoke, what the purpose of this test was and how it would fit into the broader project plan. I tried to prevent any form of performance anxiety by emphasising how it was the device, my work, that was being tested, rather than they themselves and how the whole system was still untested and potentially unstable. Lastly I asked them their permission to take an audio recording of the test and informed them that of course this would be treated confidentially and only used for this academic work.

That procedure remained the same for every test. Similar to the preliminary interviews, I would always offer the participants some form of courtesy for taking the time to help me with my work; in this case it was mostly fresh pastries and soft drinks.

The first test user was the most experienced teacher with roughly fifty years of age. She had some slight problems finding the switch to power on the device and struggled somewhat with the high buttons; they only respond when being thoroughly pressed all the way down. While I was accustomed to this, it did make the interaction for someone who was new to the device awkward. It was apparent that she had some reluctance handling the device; she expressed fear that she might break something by pushing the wrong buttons. She also seemed to assume that generally a change in state should necessitate a manual save: she searched for something like a save button both when registering new members and when changing the settings. Perhaps this could be remedied by renaming the buttons and calling them “done” instead of “back”. Furthermore there was some confusion as to where the registration screen for new members is located. Other than that, there were no major problems or hold-ups.

When asked about the display, she stated that she found the display to be a tad small, yet readable. Brightness and contrast were sufficient. She stated that she would appreciate a setting for changing the alarm sound noise level and tone. She also mentioned that she would like to have a silent mode that vibrates the device on alarm, as is custom with e.g. mobile phones. As for a touchscreen alternative, she said that she has never used touch based input on a mobile device before and thus prefers physical buttons. What she would like better about a touch based design is a more flattened surface (which was at that iteration broken by the three relatively high profiled buttons).

So the bottom line of the first interview were the following insights: the power switch would have to be made more apparent, maybe by labelling it. The menu for starting the registration of new members would have to be renamed to better reflect its functionality. To better distinguish the members list from the registration screen, a message like "No members registered yet" should be displayed if that is indeed the case. The hardware buttons need to be made more responsive and flat.

The second test user was the youngest of all four, being in her early twenties. It was obvious almost instantly how much her approach to new technology differs from that of the my first tester. Everything from powering on the device to understanding the menu navigation just happens much faster. The overall testing time was reduced from nearly thirty minutes in the first test to slightly under eighteen. She had a much franker "trial and error"-style approach than the first test user, seeming much more unfazed by the possibility of doing something wrong. On the topic of "doing something wrong": I'm hard pressed identify problematic areas with this one. Like the first test user, she initially tried the menu labelled "Mitglieder" ("Members") during the task to register new members. She also had her problems with the buttons' way of registering presses. But that's about it.

Regarding the display she stated that it was big enough and well readable, the brightness and contrast being acceptable. She regarded the noise level as being too low for practical purposes as it could be drowned out in a noisy environment. When asked about a touchscreen alternative, she said that yes, she could very well imagine the device with touch input. According to her she has had a lot of experience with touch based devices. Her age and attitude towards such input methods indicate that she would have preferred it over physical buttons, though she did not actually say so. I presume it is the case, nevertheless. That is also backed strongly by the fact that she tried to use the display like a touch-based one at first, before even taking note of the hardware buttons..

After the test and subsequent questions, she stated that she regarded the test as having been "sehr leicht", very easy. Overall I had the impression she was considering the whole thing as a challenge (yet finding little in that regard). She further remarked how the device interface had been easily understandable (which is something I am not ashamed to admit being glad to hear). Overall she was faster than test user one by almost eight minutes.

As much as I personally like seeing someone master the device's interface at fast pace, it leaves little to deduce from this test in the way of possible improvements. The only slight problems were finding the right menu for registering new members (which she found on the second try) and the haptic of the physical buttons; both also occurred during the first test. It was very educational, however, to see the stark contrast between the first two test user's approaches to a new device.

The third test user was in her late twenties, thus being slightly older than test user two. Like the latter, though, she had a very straightforward approach to handling the device. She was even a slight two minutes faster, thus being the fastest to finish all the tasks. She concluded some of them in the time it would take me to read the task description. Grasping the user interface and navigation seemed to come a bit easier overall; figuring out the functionality provided by the buttons and how to traverse the menu structures were more obvious to her than to the youngest tester. She also had a much lower ratio of successful to erroneous interactions and thus less errors to recover from. To put it shortly, she was a tad more on target than test user number two.

When asked about it she stated that the display was sufficient and well readable. She remarked that she would prefer the alarm signal to be louder. The device itself also passed muster by being neither too big nor too heavy for her liking. As for the topic of touch screens, she stated without hesitation that she would prefer the physical buttons any time. Her reasoning was that she expected them to be more robust and reliable in demanding situations (like rainy weather or high chance of dirtying) than touch based input. After the interview the tester explicitly lauded the ease of use of the device out of her own volition.

Sadly, with the good performance of the tester there is little to be gained from this as far as identification of problematic areas is concerned. The only things I noticed were that this test user indicated first searching for a hardware button (through thinking out loud) when asked to mute the device. I had not thought about this previously decided, but decided to keep it in mind for later iterations. Apart from that there was the problematic haptic of the buttons (which was about as clear a problem as could be by now) and the also known ambiguous labelling of the members list menu vs. the registration mode (which was the user's only erroneous interaction).

The fourth and final test user was of middle age, in her late thirties. Like the first test user she exhibited some restraint when first interacting with the device. The age difference to the previous two test users was apparent, yet she overcame this more quickly than the first tester. During the initial and concluding talks of the test she showed a lot of interest in the technology, which was reflected in the way she treated the device: careful yet with interest.

She took several minutes longer than test users two and three, yet was faster than the first tester. But, as I have said in the previous section, time is at best an additional metric in these tests. She actually had a very good error ratio on her actions, only missing the registration screen (the one that *every single* user missed) at first, as well as needing several tries to pass the screen lock. After the first minutes her grasp of the buttons and menu navigation was also solid. She said, on her own initiative, that she herself was surprised at her management of the devices functionality, remarking how she usually was not good with electronic devices. I fully admit to liking post-test praise.

Concerning the display, she regarded it as fine, with the readability and contrast being decent enough for her to read even on a bright day. She, too, thought that the noise level of the alarm buzzer was too low to hear in a busy environment. When asked about a possible touch based alternative, she said that, yes, she could also imagine such a device in the field, yet personally preferred the buttons for their robustness and feedback. She stated to have prior experience with touch screen usage on her phone.

This interview revealed a slightly subpar wording in the lock screen; the test user needed several tries before realising that the buttons had to be pressed at the same time to unlock the screen, instead of one after the other. A rephrasing of the onscreen description should fix that. As usual, the registration screen for new members was not found immediately, at first being confused with the members list. Also the buttons proved problematic again. I really needed to do something about the buttons. Other than that, no obvious flaws could be observed.

Summary and Conclusion

So what conclusions could I draw from these first usability tests? In short, most of the user interface worked well and appeared to be ready for field testing. There were some issues that affected all users which will have to be remedied, though.

Concerning the influence of the user's age on the performance during the test: as was to be expected, there was a noticeable correlation. Test user number three, who was in her late twenties, was quickest to grasp how the user interface worked. She did so with a noteworthy lack of interaction errors, making the right choice every time but twice during the whole test and still being the fastest to finish overall. I see the reasons for this mostly in her prior experience with electronic devices. I think it would not be far fetched to assume that she did so well with the interface because it was designed by someone that was, being of similar age, influenced by the same era of interface designs. I seem to have designed for people of my own age better than for others. The second test user, the youngest one, also performed much faster than her older peers, yet with a slightly worse ratio of successful to erroneous interactions than testers three and four. She seemed to feel more at home with touch based interfaces and less familiar with the kind of small, low resolution screens and accompanying menu approach I had opted for. Yet she quickly adapted and after a few minutes was almost up to par with test user two as far as speed and sureness of the interactions were concerned. I think that the device design, though practical, has less appeal to a younger age group. This is however speculative, as the youngest test user's performance did not differ that much from the best overall, and no additional tests were conducted to verify.

The two older test users performed in a similar fashion, though the younger of the two did somewhat better as far as accustoming to the interface was concerned. Both exhibited a reserve against electronic devices that is not restricted to my design alone. Test user four even stated such outright, claiming how she was not good with new technology. Both seemed to be afraid to break something in an irreversible fashion. They also seemed much more self-conscious than their younger colleagues. Test user four lost both reservations a bit faster, thereby performing better. Both of them had a worse interaction error ratio than their younger colleagues, but in the end managed to make do with only very few and slight hints on my part. Test user four even remarked on how she had expected it to be more complicated than that. Nevertheless, concerning the influence of age on the performance with the device I have to conclude that the older participants performed less well than their younger peers.

Some of my design choices seemed obviously flawed after the tests. The most pressing issue that had caused problems with all four of the testers were the buttons. They only registered as being pressed when pushed all the way down. This took some getting used to for all users, and even then interactions seemed strained and always involved a lot of force being exerted on the device. No doubt this would become tedious fast. The buttons also took slightly too long to register a press to be comfortable to use. This was not the hardware's fault; rather, the rate at which their state was queried seemed to be too low. These problems were obvious now, but I had overlooked them while working on the device, simply because I had become accustomed to them. With some amount of tunnel vision I had forgotten that they might be cumbersome to someone else.

Apart from these functional flaws, the buttons had some questionable aesthetics, protruding visibly from the device (Figure 3.8 on page 51 shows this in detail). Test user one remarked that she would have liked the device to be more flat on the front side, naming this one of the most prominent advantages a touchscreen would have. She is of course right, for practical reasons (e.g. when carrying the device around in a bag or pocket), the buttons should not stick out that much from the rest of the device. It seemed pretty much unavoidable by then to completely replace the hardware buttons. I would also have to think about making the power switch more obvious, since it was not always apparent as such immediately.

There was also one erroneous interaction that all participants had in common, which was to look for the registration screen at the wrong place, namely the option that led to the list of currently known members. As this had happened in every test, it was most likely a fault in my wording or the menu structure. I decided to do a restructuring of the member management menus, or at the very least change the menu descriptions to be less ambiguous. Furthermore I would need to change the wording on the lock screen, as two of the four testers had initially tried to press the buttons one after the other, when they should have done so simultaneously.

Furthermore, three users had seemed doubtful what to do after having registered devices or changed an option. They made the impression of being unsure whether or not current changes would be lost if they went back to the menu they came from. I had thought that immediately saving changes without prompting the user to save would ease the usage, e.g. if a setting is changed, this is made persistent without further user input. Indeed the opposite seems to be the case. The participants seemed to expect a "save" button, fearing to lose data otherwise. I resolved to rename the buttons from "Zurück" (back) to "Speichern" (save) or something similar.

Lastly all of the participants said that they wished for an option to somehow influence the alarm signal. These requests ranged from control over the noise level to changing the tone of the alarm. I resolved to look into the matter, but had little hopes concerning the sound level. The piezzo buzzer I had used in the prototype only allows for a variation of the sound frequency, yet not the noise level.

Apart from those listed above, I was unable to identify further problems with the device interface. The overall reception of the prototype was positive. Three out of four participants remarked that they had gotten along very well with the user interface and lauded the ease of use. This happened without incentive on my part (for those readers that are suspecting I go around asking people to say how much they like my inventions) on the user's own accord. I saw this as a very good sign as far as device interactions are concerned.

The alarm signal worked well and was instantly recognisable as such to all users. I threw the alarm into the test without any advance warning so that the user would be confronted with it as unexpectedly as possible. All testers correctly interpreted and acknowledged the alarm and message.

The device and interface seem to be simple enough to learn without manual or other aids, even though all the necessary functionality for a field test was already there. I only had to offer some slight help to test user one. All other participants made do without assistance. All tester came to grips with button usage and menu structure. I did not so much as explain a single button before the test, they all managed to figure it out by themselves.

The physical aspects of the device were also received well enough. When asked, the size and weight were acceptable to all the participants. The buttons were easy to reach for all of them. At the beginning, all users operated the device with one hand, using the thumb to press buttons as can be seen on Figure 4.3a. Over the course of the tests, however, the preferred way became to hold the device in one hand and operate the buttons with the other. Figure 4.3b shows an example of this. All users were able to read the display without problems (or glasses).

(a) operating buttons and holding the device with the same hand

(b) using one hand to hold the device and the other for operating buttons

Figure 4.3: Two different examples of how the main device was handled

The buttons, even though they had their difficulties, were preferred by three of the testers over a touchscreen implementation. Their reasons for this choice were the same as mine, namely robustness and reliability. I have to admit, however, that there are caveats to the way I posed my question. I did not provide the user with an example implementation of a touchscreen based device. I merely asked them whether or not they would have preferred such, or could imagine it as an alternative. Their answers would have been much more significant if I had actually had a second, touch based device to compare. Since there were constraints on time and money (as there always are), I did not happen to have such a thing. Thus, all I can say with certainty is that the buttons work, not whether they are better than the alternative.

Lastly there are some general observations to be made that are not directly related to the performance of the device during the tests. For setting the range of the device, I had opted to use fuzzy terms such as “near”, “intermediate” and “far”, refraining from making promises like “50 metres” or such. Interestingly, none of the testers was interested in how far “far” or “near” actually were. There were no questions concerning that; this observation aligns with some remarks from the preliminary interviews about how such concrete indication of range was of little use (see page 18). Then there was the observation that only one of the testers (number three) used all three of the device’s buttons. The others just utilised the centre button and either the left one or the right one as soon as they discovered their usage. So in the menus, they just moved the selection in one direction and if they missed the target pressed the button until they had traversed the menu once and were at the beginning again. None save one user thought of trying the other button to move the cursor in the opposite direction. I was unsure about the implications of this. Should I really get rid of one of the buttons as it had not been used? I decided against that, since I saw no gain from reducing the number of buttons; an expert user may well start using all of them. However, I have none such users available to verify that. Thus the third button received the benefit of the doubt.

Conclusion As to the goal of this usability tests, to see whether or not the device was ready for field testing, I concluded that this was indeed the case. The device was usable and had been received well. The display was readable, the alarm noticeable, the interface understandable and the device had been stable throughout all of the tests. I would make some slight improvements to the hardware and do some rewording on the software, but otherwise I saw no reason to not put the prototype to the real testing.

Reference You may find a documentation of all the actual changes to the prototype, which resulted from this usability testing phase, starting on page 57.

4.2 First Field Test

It is stated in [81] that the fidelity and aim of an evaluation should dictate the choice of method, between an examination under laboratory conditions or in the field. A low fidelity prototype should not be tested in the field; an evaluation under laboratory conditions would be better suited. Yet to discover the usefulness of a device, a field study is highly recommended, as it allows to better discover problems arising from mobile use. Thus I felt that, after the initial usability examination, my prototype was ready to be tested in the field, so that I may discover aspects that would have escaped me otherwise.

The first field test was undertaken on June 6^{th}, 2013, starting at 13:00 and ending at 16:00. Goals of this endeavour were to test the acceptance of the devices by children as well as the technical functionality in a real world environment. I wanted to see how children would react to the items I had built; of specific interest was whether or not they would wear them without complaining, if they would feel hampered by wearing them or if it would change their normal behaviour somehow. The second aspect to test was of course whether or not the system would work in actual real world scenarios. Up to that point I had only tested the devices in laboratory environments; movement and the full number of devices was used, yet an actual field trip with a lot of children turns out to be pretty much impossible to simulate. So this would be the first time that any of my work would be submitted to "real" use.

On the Importance of Field Testing

I figured field studies to be essential when trying to evaluate a mobile device. As is noted in [81], usage of mobile devices involves much more than becomes visible under laboratory conditions. The use context for the device is more intense and changing frequently, from a walk in the park to crossing a street to taking a public transport. Hence problems will be revealed that were not obvious by observing the user in a calm and controlled environment. In [72], Nielsen *et al.* take a strong stance in support of field testing. They conclude that it helps to identify significantly more usability problems, which in their opinion is well worth the added complexity and expense of time. These findings are shared by Duh *et al.* in [43], where they state that many usability problems were found only in the field and how those tend to be critical ones. They, too, state the environmental factors like noise or added cognitive load as a reason for this and note how such factors cannot be ignored if one wants to conduct an accurate usability evaluation.

Preparations

A week before embarking on the first field test I composed a letter that was to be given to all the children, which they were to take home to their parents. In this letter, I shortly introduced myself and explained how I was to accompany the children on their excursion as part of my master thesis. The text described the project that I was working on and why it would help the teachers. With this writ I asked three things of the parents: that they would allow me to offer the items to the children for wear, that I would be allowed to take photos during the trip and that I may ask their kids a few short questions. I assured them that all of this was voluntary and in no way connected to mandatory school tasks or grades. I also included my phone number and mail address as well as that of my supervisor in case any of the parents should have further questions (there were none).

At the end of the letter there was a section for the parents to cut off, sign and return; this is the usual way of day-to-day school related communication between teachers and parents. 22 parents granted me these permissions. Three children were explicitly not permitted to partake - they were neither included in any photos nor asked if he wanted to wear a device. The parents of one other child did not understand the letter and asked for a translation into the Turkish language.

Before the beginning of the field trip the children were gathered in their classroom. Their teacher told them who I was and why I had come here today. I showed them all the items I had brought with me and shortly explained what they were meant to do - that they would help their teacher to keep an eye on them and prevent anyone getting lost. I also told them that I would be accompanying them today to see if my "invention" works.

They seemed to understand and accept that well enough and had a lot of questions about the items and my work. General mood was excitement (about a new person coming with them) and curiosity (about the devices). All the items were laid out on a table and the teacher asked the children who wanted to carry which device. All who had permissions from their parents immediately volunteered; five were picked by the teacher with the others being promised that they would get their chance to also try an item later. After the children had their items equipped they were registered with the main device.

It was observable at this point that the registration mode needed improvements. First of all the timeout for locking the screen was still active. This resulted in buttons having to be pressed regularly to prevent the device from going into standby. That was a design flaw which I would remedy until the next field test. Secondly, the child devices were programmed such that they would only attempt registration once, about five seconds after being switched on. This resulted in some rush during registration, because all devices had to be powered on just then. This is something that I would also change until the second field test.

There were 25 children present for the excursion. They were overseen by their class teacher as well as one "Freizeitbetreuer". The latter's job is to supervise the children during the off hours when they are not in class, such as during all-day care. So this person was not a teacher, but nevertheless accustomed to watching over the children. After gathering outside their classroom, the children left the school building in an orderly fashion. They were joined by another class of the same size; they, too, were overseen by their class teacher and one additional adult supervisor. Both classes would undertake the excursion jointly, resulting in 50 children who were watched over by four adults.

Transit

While walking from the entrance of the school building to the bus, the children proceeded in an orderly fashion, walking in a two wide column. One teacher was leading the front, while another was bringing up the rear. The two additional adult supervisors were walking alongside the children, between them and the street. This way, the children were keeping close to the buildings, leaving ample room between themselves and the streets. They remained in this formation while waiting for the bus to arrive at the station. During this phase the device worked well. There was one false alarm, which was due to one of the child devices failing. I suspect a short circuit with the conductive thread, but the problem did not resurface and I could not reproduce it.

As soon as the bus arrived, the children started pouring in through two of the four access doors. The teachers were waiting outside and keeping overwatch to assure that everyone had entered the bus. Only once they were sure that all the children had boarded did they enter themselves. Meanwhile the children had spread out within the vehicle so that the group stretched from back to front, interspersed with regular commuters. The bus was of the large, low-floor articulated type that is used in the Viennese public transport system (MAN NG 273), which is approximately 18 metres long. The children wearing my items were about evenly distributed over the whole passenger area, while the main device was located in the back. During the ride there were several false alarms, rendering the system unusable for all practical purposes. I attribute this to the fact that there were a lot of interferences from the surroundings (lots of metal and people) inside the vehicle. Furthermore the low number of child devices would prevent effective mesh networking over the whole length of the transport. Upon reaching their destination, the teachers told the children in a loud voice (yet without shouting) to exit the bus. Again, they stayed behind until satisfied that all children had exited the vehicle. This proved a bit more hectic than boarding, since there were a lot of other people present and many corners where one child could remain behind unnoticed. They seemed to mostly rely on the children passing around the word that this is the right stop to get out.

(a) close formation

(b) loose formation

Figure 4.4: School class walking through the park

Once inside the park, the teachers decided that the group should first walk around a bit. They let the kids run in a loose group with at least one adult walking in front and one bringing up the rear. No child was allowed to walk behind the last supervisor. The constellation of this procession was constantly changing, as kids were moving from the front of the group to the back and vice versa. Figure 4.4 depicts this situation, with 4.4a showing a close formation, which was mainly observable early on and soon dissolved into something more resembling what may be seen in Figure 4.4b. Occasionally, the teachers called a stop at key points along the route, such as lakes or memorials (situation is depicted in Figure 4.5); after a short duration the group would move on again. I would estimate the distance from the first child to the last to have been about eighty metres, give or take a few. While walking the teachers would spend a lot of time talking to and entertaining the children. They seldom had time to actively check the device. This means that a well audible alarm signal is a must have. The signal as it is was heard, but barely so. Making it louder would not be wrong. More so considering that the device was held in hand for the whole

Figure 4.5: School class making a stop during their walk through the park

duration and not stored in a pocket or bag.

Walking around the park took about an hour. Throughout this phase the device worked well with only one false alarm being given. When taking into account that the way the children spread out within the group was completely random and only five devices being present to form a network, I consider this to be a good result.

Playground

After walking around for some time, the whole group gathered at a playground within the park; Figure 4.6 shows this location. The children were allowed to roam and play as they pleased, so long as they did not leave this enclosed area. At that point, the children who had up to now carried the prototypes were told to give them to others that had not previously had a chance to try them. This was actually necessary because most of them were loath to let go of the devices and there was some demand to try them. After that the children split up into small groups, some sitting around on benches and some running and playing, trying slides and climbing around – activities that children tend to do when on a playground. The adults were talking among themselves and to some of the children that were not playing around. They did not seem to actively supervise the children - sure enough they always had one eye on some of the “hot spots”, but also seemed to rely on the children obeying their directive to not leave the grounds. It was stated on several occasions by different people during the preliminary

Figure 4.6: Playground visited during the first field test

interviews that this is a situation where a device such as the one that I developed would come in handy. And indeed this was were, out of all the scenarios encountered during this trip, the device performed best. There were no false alarms, no hiccups and no outages. I attribute this mostly to the semi-stationary nature of this phase.

Observing the children I tried to gauge any change or difference in their behaviour and movement, as opposed to that of their friends who did not carry one of the prototypes. I am pleased to say that I did not discover any. Those kids that wore one of the devices did not behave any different with or without, I observed after they had switched around the items within the group once more. They ran around and played unhindered and unconcerned with the additional apparel. After being initially quite proud to wear it, they soon seemed to forget that they were even carrying some foreign objects.

During the stay I would also explicitly ask children wearing the device how they liked it, in the event that they happened by me and were not focused too deeply on their play. All of them stated that they liked the items and that they did not feel inhibited wearing them at all. This coincides well with my earlier observations. Some items were especially liked for their appearance: notably the grey cap was liked by the girls while the boys were into the high visibility vests. I quote one boy saying to another: “Ur geil schaust du aus”, which is very colloquial and roughly translates to “You look rad”.

Figure 4.7: Children wearing prototype devices during the field trip. Items from left to right: brown cap, high visibility vest (twice), grey cap and pendant

The group remained on the playground for about an hour. When gathering the children, the teachers went around the area and told everyone that they were now preparing to leave and it was time to gather at the exit. This process did not seem to be rushed or time critical, nevertheless I think the notification function of my device could make it less arduous. Sadly, there was some confusion and I did not get to put this functionality to the test. I resolved to make up for that negligence on the second field trip.

Once the children were gathered, I asked around how they had liked the items. They all agreed loudly that they had liked them well. Next I asked them what they had liked best. The high visibility vests came out on top quite clearly in this comparison. Most said that they had liked the vests best both for the appearance and the wearing. Next came the grey cap, being mentioned by about half of the girls present. Neither the brown cap nor the pendant were mentioned, either in a positive or negative way. The teachers stated that they liked the vests a lot for their high visibility.

A note on the durability of the devices: it was rainy throughout the whole day, with intermediate stronger showers. The devices were in use throughout these and did not show any sign of malfunction or outages. They proved to withstand a certain degree of splash water. I am less optimistic about extended periods of being exposed to heavy rain and will refrain from testing

Figure 4.8: Two children wearing the high visibility vests

them in such conditions.

Summary of the first field test

In summary there were three distinct scenarios in which my system was put to the test. First there was the transit within an urban area. Here the devices exhibited some severe lack of reliability. While the walk from and to the bus worked without incident (spare one device outage), the system was rendered unusable inside the vehicle. In the second phase the whole group took a walking tour through the park. There was a lot of individual, random movement within the group, which was spread out more than in phase one. Nevertheless the system worked far more reliable here. This would be due mostly to the relatively calm surroundings (as far as radio interferences are concerned). There was not even much necessity for mesh networking, since the devices could cover the distance from the back of the group to front in one hop. The third and last phase was the stay on the playground area. This is where the devices worked best, since the environment was almost optimal. There were little to no obstacles and the movement of the children was confined to an area small enough to allow all items to remain in contact with the main device.

Conclusions

- The devices were well liked and received by the children. How this holds up once the novelty factor wears off will have to remain subject to speculation for now.
- High visibility vests would be my item of choice if I wanted to advance this system beyond a prototype. They were liked best out of all the devices, both by the children and by the teachers. Additionally they proved to be the least prone to errors and provided good radio signal quality.
- The systems shows severe vulnerability to interferences. This was, however, to be expected due to the low number of devices as well as the low power of the transmitters.
- The bell of the main device was heard by the teacher, but only when the device was held in hand. In a more noisy environment with the device being put in a bag or pocket, this will not be enough.
- Five devices are seemingly not enough to do efficient mesh networking. I will consider configuring them as end devices and discarding multi-hop networks, using signal strength for distance measuring.
- The registration of devices needs to be made more stable and less prone to becoming hectic.

I cannot help but remark that trying to bring structure into a field test with a group of young children is a futile endeavour when being undertaken by one single person. I had prepared a list of things to observe, a list of questions to ask, a list of tests to undertake with the main device, a list of things to document through pictures – all of which was useless in the end and did not come to bear, as far as providing structure to the observations is concerned. Children of that age exhibit such a remarkable curiosity, an inquisitive nosiness about everything new to their environment that was, from start to finish, constantly wrenching away my attention from plans I had made. There was no way to be an unobtrusive, outside observer in this - they simply would not let me remain "outside".

My plan for the next field test was this: I would try catch up on testing the notification feature of the devices and observe the teachers handling of the device more closely. Until then I would try to improve the system stability of the caps as there were some system outages with both.

4.3 Second Field Test

There were several things that I wanted to observe during this second field test. Like with the first venture, I wanted to see how the children put up with wearing the devices. Specifically, I was interested in whether their attitude had changed now that the novelty factor of the items had decreased somewhat. Furthermore, I wanted to individually interview those children that had been given one of the items. In the first field trip I had only talked to the whole group. Besides that I wanted to try and focus more on how the teacher handled the device, something I had neglected somewhat outside of laboratory environments. Lastly there was the feature of signalling the children, which still remained to be tested.

Originally, I also wanted to field test a new networking approach where all child devices where configured as ZigBee end devices; this was also meant to feature distance estimation via signal strength readings. Sadly I could not get that setup to work at a satisfactory level before the date of this field test, so I had to cancel it. Please refer to 3.3 for more information on this effort. This failed attempt left very little time for other endeavours as far as device improvements are concerned. What I did change was the registration mode for adding new child devices to the group. Firstly, I disabled the screen lock timer while in registration mode. Formerly, the user had to press a button at regular intervals to prevent the device from locking (and thus leaving registration mode). Secondly, I reprogrammed the child devices such that they would communicate better with the main controller and attempt registration each time they received an according response code to their reports.

Preparations

The class I would accompany was the same as during the first field test. This time there would be two teachers to supervise 22 children. I gave the class teacher a short walk-trough of the main device to freshen up her memory about its functionality. There would be another class joining us on this field trip (supervised by two other teachers), as was the case also last time. From observations made later during this trip as well as the first one it would seem that, though each pair of teachers is responsible for their respective class, they tend to manage the group as a whole during transit.

I tried to keep more to the background this time while the teacher distributed the items to selected children. As the children were already familiar with the project, they did not need any further introduction – they seemed to remember the purpose of the devices quite well. Interestingly, their enthusiasm for obtaining one of the prototypes did not seem to be diminished much, which was contrary to what I had anticipated. I hypothesise that this

Figure 4.9: Teacher distributing items at the start of the field trip

is due to the scarcity of the wearables: there were only five of them to go around. I think it would not be unreasonable to assume that this shortage has no small part in the children's sustained desire be the wearer of one of the devices. However that is all speculative and the children's long term acceptance of extended deployment of such devices shall not be the subject of this work.

Transit

Transit procedures were the same as during the first trip. The two classes met in front of the school building and then jointly proceeded to the nearest tramway stations. The children formed a queue and one teacher took the lead while another one brought up the rear. The two remaining teachers usually positioned themselves halfway down the queue and tried to keep the children away from the street, towards the building side of the pavement. When we had to cross streets, one of the teachers would always stand in the middle of the road to enable the children a save traversal. At larger junctions with regulating traffic lights the procedure remained largely the same with an added time constraint (due to traffic lights changing) and the resulting fuss to get all the children across as fast and yet orderly as possible. At this point, I first noticed that the teacher completely ignored the device she held in her hand, never so much as glancing at it. Instead she relied firmly on her best practices.

Figure 4.10: Crossing a street

When entering the tramway, the teachers first let the children board by several entrances stretched over most of the train. Only when they were sure that none of the children remained outside did they enter themselves. Inside the train children settled were they found place, scattering over almost half the length of the tramway. During the ride the teachers seemed to be well aware of where their children were and what they were doing. There was no counting of present members though, most likely due to the wide scattering of children throughout the train. It is easy to see why teachers listed these transit scenarios as the most crucial as far as keeping track of the children was concerned. Both during the bus and train ride there was little order to be had in the movement and dispersion of the children once they had boarded the vehicle.

Exiting the train was even more of a challenge. Before boarding, the children had waited in what goes for rank and file at that age, making it easier for their supervisors to see if all had gotten aboard the train. While exiting though, there no order to be found. To me, someone with little experience as a supervisor, it seemed like there was no real way to keep overview of all the children dispersed throughout the train cab. Once outside and after some order had returned to the group, the class teachers started counting the children; you can see this procedure taking place in Figure 4.12. By the time they had finished, the train was long gone from the station. I can very well see how teachers would appreciate some support in these situations.

Once the counting was done and every child was accounted for, the group started moving towards the park. This took about ten minutes on foot. The procedure was the same, except now the queue was stretched thinner, due to the more narrow pavements. There was little to no communication between the teachers when moving in this order. They were too far from each other to talk normally and seemed occupied with watching the children anyway.

Figure 4.11: Waiting to board a train

During the whole transit phase, the devices exhibited the behaviour I had come to expect from interference-rich surroundings: there was one false alarm on the way from the school to the tramway. Inside the train there were multiple more. Sadly, I have to conclude that this prototype is just not ready for such an environment, and will not be within the course of this work. Whether that could be remedied by adding more child devices to the network to better utilise mesh networking remains yet to be seen.

The Park

Once we had reached the park, the group made for a nearby playground – a twenty minutes walk away from the entrance. As the last time, the teachers would let the children disband from their usual queue formation. They were free to run about so long as they did not stray too far from the main body of the group. At least one teacher would always bring up the rear and make sure that there was no child left behind. Sometimes the supervisors would intentionally let the children run on ahead, giving them a “target” to run for. As was with the last trip, the device managed fairly well in this situation, giving one false alarm that was due to the child in question being out of range; an estimated 80 metres away from the teacher carrying the device. Figure 4.13 illustrates this situation.

Figure 4.12: Counting of heads after exiting a tramway

The playground was of larger proportions than the one we visited on the first trip. The children were allowed to play and linger anywhere in the area, including a small stream that ran alongside parts of it. The teachers set themselves up at the edge of the yard, about 30 metres away from the actual playground, where most of the children stayed. They occasionally did a round through the area, but mainly stayed and talked among themselves. Due to this constellation, with the teachers being on one end of the area and the children being on the other, the average distance between the devices was greater than during the first field test. The greater distance meant that the devices did not work as flawlessly, with several alarms while the kids moved in and out of range. Positioning the device more closely to the bulk of the group remedied this problem.

Observing the children scurrying around on the play yard I noticed that I had underestimated how much the human body absorbs radio signals. The high visibility vest and pendant devices would sometimes fail to report with the main device for several seconds if the children had turned their backs on it. Once they turned around again, the reception went back to normal. The brown cap, which had had the transmitter integrated into its visor, showed this problem to a lesser extent. The grey cap, with the transmitter situated slightly above the child's ear, proved to be least susceptible to the position of the wearer relative to the main device. The blockage of radio signals by the body is a design flaw with the devices that should be remedied in any further implementations.

Figure 4.13: The class on its way through a park

Figure 4.14: Playground visited during the second field trip

During this stay on the playground, the wear and tear that the child devices will have to endure became more obvious. All of the devices received some level of dirtying. While the children played on slides, swings and around (as well as once, unintentionally, in) the small stream next to the yard, they showed little care about what kind of dirt that would leave on all of their clothing. That is to be considered typical playground behaviour at this age. The high visibility vests sustained the most visible traces of wear: They were muddied and partly wet mostly in the lower half of the front and all over the back. Figure 4.15 shows one high visibility vests in the state it was in after both field tests. The pendant received some slight dirt all over, as can be seen in Figure 4.16. All three devices were in need of a washing after the trip was done. This shows that it is important for the devices to not only be robust, but also easily and regularly cleanable. Perhaps it would be better still to fashion the electronics such that the clothing could easily

Figure 4.15: Wear and tear visible on one of the vests

be interchanged in case it becomes smudged beyond washing.

The caps did not get dirty, yet as our stay on the playground progressed farther I saw another flaw become more relevant than it had been up until now. Due to the warmer weather and the children's increased physical activity while playing, the caps had become all but soaked in sweat. Consequently, after about an hour of our stay the children did not want to wear them any longer. This happened even though temperatures were far from hot for Central European conditions; weather stations reported around 27°C with little to no wind. That illustrates the limited field of application that headwear has for the purpose of this system: caps are reasonable to wear for temperate weather conditions while they are too hot for even slightly warm temperatures when paired with extended physical activity. They will be too cold in late autumn and winter.

On a side note, one teacher's forecast about how the caps are too easy to take off and therefore prone to being discarded by the children (see page 18) proved to be quite accurate. As soon as the children did not want to wear the headgear any longer, they silently took the items off and placed them on the bench were I was taking notes. I didn't even notice this on

Figure 4.16: Wear and tear visible on the pendant

two occasions, just finding the devices sitting next to me, after I had been busy writing. Such behaviour was not observable with either the pendant nor with the high visibility vests.

An irksome fault with the design of the high visibility vests did also surface as a result of the wear and tear during field tests: on both items, the connections on the battery pack became loose, resulting in a total loss of power and and two useless devices. I tried my best to make the wiring of the energy supply as sturdy as possible (it can be seen in Figure 3.20), yet the strain seems to have been too high; the wiring became loose during the stay on the yard, presumably during play, and broke directly at the battery pack. In my defence, I have to say that the connections are pretty frail ex-factory; the wires are only held in place by a small blob of solder. I broke two of these, one being the bigger battery pack implemented in the main device, during my work on the prototypes.

Something that I found interesting about the behaviour of the children was just how unconcerned they were with the devices. While I am happy to see their high acceptance, in retrospect it baffles me how little attention the items received once acquired by a wearer. Some children noticed the lights shining under the textiles, but that was as far as their curiosity went. They almost seemed to forget that they were wearing unfamiliar equipment. One boy became very inquisitive about the range of the transmitters, so to quench his curiosity we did some coverage tests. However, none of his classmates showed any similar interests. Their unconditional as-is acceptance of the devices is unexpected, but not necessarily a bad thing.

Interviews

While the group stayed on the playground, I took the time to interview some of the children individually. Their teacher helped me pick the ones that had carried at least one device. The interviews were done on a bench at the edge of the area, a little detached from the rest of the children. Far enough to reduce distractions by others, yet not so far away that the child would feel uncomfortable. Additionally, the class teacher remained with us for the whole duration of each interview to help make the children feel at ease. She would also encourage them from time to time to talk freely and help clarify once or twice when I had phrased a question badly; this was very welcome.

The interviews were very short, consisting only of three to four questions asked of each child. The children were eager to get back on the playground. As such the interview conditions may have been below optimal. Still, the children felt well at ease in these surroundings and the devices were still fresh in their memory, so I decided to go through with the interviews anyway. The talks were done in a semi-structured manner. There were three main questions which I posed to each child, but let them talk freely otherwise. The following topics I tried to address in each interview: how did they like the device they were carrying? If they had carried multiple devices, which one did they like best? Do they feel uncomfortable if their teachers always know whether they are near or far? Is there any other kind of clothing that they would rather wear? Would they be OK with carrying such a device on every school outing?

Overall I talked to eight of the children. On the topic of which of the devices they liked best, the high visibility vests were chosen by six of them. The other two chose the pendant, mostly for it's visual aesthetics. On that note, the children seemed to generally choose what they considered to be pretty and mostly disregarded whether or not that item would be convenient. Only one child stated that it liked the high visibility vest best because it was comfortable to wear. Interesting at this point is that none of the children favoured the headwear items, even though three of them had worn those. The caps seemed to have fallen from grace. The children all agreed on the topic of how they felt about their teacher always knowing whether or not they were present. None of them seemed to find the notion unsettling or uncomfortable. They all stated thus without hesitation or afterthought. It would appear that this topic was of no real concern to them. Interestingly enough, even these young children understood perfectly well the importance that their teachers always know of their presence so they would not get lost.

Another thing they all agreed on was that it would be acceptable to wear such devices as they had carried now on every school trip. When it came to ideas for further improvements, two children brought up items of their own: one said that he would like something to wear on the wrist, similar to a watch. The other one said that he would prefer shirts, which the school had issued to them, featuring the school's emblem and name on the front. The first idea had been on my list for possible prototypes, but had been discarded because any Arduino components I could find were ultimately too big and clunky to be worn on a child's wrist; otherwise the idea still holds promise. The second one lacks some feasibility for varying temperatures and weather conditions: while fine in summer, they would have to be worn under at least two other layers of clothing during winter. Additionally, as the first layer of clothing a shirt could not be worn above any other personal clothing, like the high visibility vests. It would be interesting to find out whether the school's logo and name on the items would give the children additional incentive to wear them. However, this question will have to wait until another time.

As a concluding remark on these talks, I have to say that I found it somewhat difficult to effectively interview the children. They mostly responded with single sentence answers to my questions, hardly ever doing any free talking of their own. This did not change whether or not their class teacher was present. Maybe they were uncomfortable because they had only seen me once before; though if that was the case, they did not show it. They also seemed perfectly at ease with me during both trips otherwise, talking to me and riddling me with questions. They did not seem uneasy, just not very talkative. That is something I would not expect from children at that age. I was a bit disappointed with how little feedback I managed to get.

Usability Assessment in the field

It is difficult to really assess the teacher's handling of the device in an actual real world environment. Every habit they show in hectic situations they have acquired through experience; those are tried and tested measures and they work. This is problematic insofar as the device is foreign matter to it all. The teacher in question used the system during two field trips, yet has year and years of experience making do without such a device. Obviously her habits prevailed in situations where my system was meant to help, resulting in her forgetting about it and simply carrying it along. I did not want to force her to use it either, since that felt like intruding on well founded best practices. Perhaps that is some valuable information gained from these field tests: when things get rough and hectic, most teachers have a set of habits, gained from experience, that help them keep their charges safe. During these situations a device such as mine could prove a distraction for experienced teachers rather than the help it is intended to be. It would be prudent to take into account that there are a lot of varying best practices already in place. These should be respected and if possible supported rather than intruded upon. Let the teachers decide when and if at all they want support.

There are two caveats with me trying to assess the device the way I did. Firstly the teacher only came into contact with the system three times, once during usability testing under laboratory conditions and twice during field tests. As I have stated at length, experienced teachers have useful best practices which they fall back on. Because this was a real world setting (and presumably that is what field testing is all about) with real world dangers, the teacher put the device to the back of her mind once her experience and routines came into play. I can hardly urge someone to better utilise my prototype at the expense of a child's safety; that is one of the rare cases where "for science" is not an acceptable excuse. It would be necessary to conduct a long term observation to see if, in time, a teacher would incorporate a system such as mine into his or her habits. That is beyond the scope of this work. The second caveat is that the system, as it was tested during the field trip, was incomplete. It comprised of only five child devices. Thus the teacher had to do both the usual checking and counting of children, no matter whether or not she wanted to utilise the prototype. Consequently, she frequently forgot about using it.

On a positive note, I think that my general interaction scheme worked well. The device was passive and unobtrusive, letting the teacher do her routines without interruptions. So unobtrusive, in fact, that she sometimes seemed to completely forget it. On the occasions when an alarm was given, the device was still able to call attention to itself and quickly inform the carrier who was missing.

Summary of the second field test

In conclusion, this second field test yielded the following insights: The high visibility vests again proved to be the most apt kind of wearable item for my purposes. Their visual appearance appealed to the children and they allowed for unrestricted movement on the playground. They also exhibit good characteristics as far as practical purposes are concerned: the vests may be worn above other apparel that the children come to school with. Being the top layer of clothing, they could also provide good radio signal quality. However, due to a design flaw on my part, this is currently only partially implemented. The transmitters being placed on the chest area make the signal strength very dependant on which way the child is facing from the main device. The field test has shown that the radio transmitters would have to be relocated to a more exposed body region in further iterations of the device, possibly to the shoulders. A second fault with the engineering of the high visibility vests concerns the battery; as noted on page 105, the connections broke on both vests.

Figure 4.17: Teacher counting children before leaving the playground

The headwear items, on the other hand, proved to be the least viable choice. Their field of application is severely limited by seasonal weather conditions. They proved to be too warm for extended wear even on a typical mild summers day in Central Europe. On the other hand they will not be warm enough during any time between late fall and spring. After about two hours of wear they quickly fell from grace with the children. The pendant is somewhere in the middle of the two, as it is easy to attire and appealed to the children visually, but suffers the same radio signal strength limitations as the high visibility vests.

As mentioned at the beginning of this section, there was the issue of testing the notification on the child devices; as you may remember, three of the five devices had this feature implemented. The two caps were equipped with LED lights at the front of the visor and one of the high visibility vests featured a buzzer on the chest area for auditive signals. Unfortunately, this plan fell through. I wanted to test it first when leaving the playground as a signal for gathering and then again on the way back in the tramway as notification to get off. What foiled this intention was, as I have noted earlier, that a) the children did not want to wear the caps any more because they were practically soaked in sweat and b) that the high visibility vest's battery died. Hence no field testing was done with the notification feature on the child devices.

Interviewing eight of the children revealed a high acceptance of the devices and their purpose. I was under the impression that all of the children understood and acknowledged the need for their supervisors to know of their presence or absence. They did not express any concerns or doubts when confronted with the possibility that their teachers be always aware of whether or not they were too far away. This is backed by observations on both field trips: the children wearing any of the items did not behave in any way different than their classmates, there was no sign of added restraint. Neither did they seem daunted at the prospect of having to wear uniform apparel at each school outing. The majority (six out of eight) chose the high visibility vests when asked which item they preferred.

As for the field assessment of the main device's usability, I found that difficult to conduct. The reasons for this are twofold: the teacher had to incorporate a device that she had only used twice before into routines and practices she had acquired through years of experience. That did not work too well, insofar as she often forgot that she was carrying the prototype. Secondly, the system was not complete, covering only about 20 percent of her class; she was forced to do the traditional counting of heads at key points anyway. On the plus side, I observed that my device was unobtrusive enough to be "forgotten", i.e. be carried around without necessitating usage, yet still was able to call attention to itself in case of alarms.

Chapter 5

Conclusion

To conclude my work, in this last chapter I will give a summary of what was achieved and what may be inferred from these results. I will also, for the sake of enabling any reader to do a true assessment of my results, provide a list of caveats that I see in my own work, which may have affected its outcome. Furthermore I wish to show some possible opportunities for future work in the event that anyone deems this topic interesting enough to devote some further research to it; I would certainly hope so, as I personally think it holds promise.

5.1 Conclusio

In this work, I have examined the possibility of aiding group management by assisting with awareness of physical presence. I have introduced a problem statement about how it is often difficult for a single person to keep an overview over all members of a group and check for absentees, specifically choosing the scenario of a school class on a field trip as primary incentive. Preliminary research, which was done extensively, has yielded no prior work done on this particular topic; however, much has been written about related fields, such as wireless sensor networks, human interface design for mobile devices, range and position estimation as well as wireless communication technologies. Many of these works were taken into account if their findings proved relevant. I have, at the beginning of this thesis, proposed a solution that facilitates some mitigation of the difficulties posed by the management of mobile groups: An electronic assistance system tasked with checking the presence of group members and giving notifications about absentees. Through this, I hope to ease the cognitive load of supervisory personnel. I have led preliminary interviews to gain a better understanding of the target audience and the problems they face, which also served as a requirement analysis. I have provided theoretical groundwork by assessing currently available wireless technologies and providing a founded choice of best fit. I have gone into the topic of the theoretical aspects of range and position estimation, especially in wireless sensor networks. Furthermore, I have provided an implementation of the proposed solution, using the Arduino prototyping platform to serve as a proof of concept that the idea is viable. This prototype was described at length, featuring one main device and three different types of items to be carried by children. Its development has been an iterative process, taking into account much input from the target audience and using [73] as a guideline for the most part. Evaluation has been done through usability tests under laboratory conditions as well as in the field. Both test methods have yielded valuable input as to the design and usage of the device. I will, in this section, give an analysis and discussion of the results of my thesis in the light of topical prior work and academic literature.

With all that being said, what are the results of this whole endeavour? The preliminary interviews have yielded, besides a good insight into the teacher's work, three types of scenarios that are most critical during a school outing. These were: Entering and exiting a public transport vehicle, dissipation of the group over a wide area and the passage through highly frequented public spaces. According to the teachers, all of these situations pose an above average risk of a group member getting separated from the group and consequently being lost. Each of these was mentioned by multiple participants, showing that their opinions of problematic situations overlapped greatly. As I have myself observed these scenarios during field testing I can only agree with their assessment. The problem scenarios are interesting and in my opinion transcend this specific use-case, in that all of them pose a different set of challenges for an implementation and help to identify requirements. The boarding and exiting of a public transport vehicle necessitates a close maximum range to quickly detect stragglers being left behind, while staying in a wide open area requires a wide range. Both scenarios also benefit from giving the teacher a possibility to notify children of an important event, for example the right stop to get off or to gather at an agreed location. The third problem scenario needs a medium range setting, yet has a high probability of showing heavy interference. All three of them possibly occurring on the same field trip reveals the requirement of having the option to choose from different range settings.

In my assessment of wireless networking technologies, I have made a point for using ZigBee. The rationale behind this choice was that it offers low energy consumption, mesh networking and is widely available as well as relatively cheap. After having worked with this technology to implement the prototype, I stand by that decision. ZigBee offers good scalability, which is an important factor due to the variable group size, through its mesh networking infrastructure and proved reliable enough during field testing, especially in open terrain. As for urban environments, I hope that mesh networking with more devices than were available during the field test will mitigate the high interferences there. However, this remains to be seen.

The question of whether or not my proposed solution, aiding in group management through assisting with awareness of physical presence, is actually implementable was one of the most important ones to be answered in this project. In conclusion, I can answer it in a positive way, albeit not without flaws and deductions. The basic functionality, to check for absentees and display their names to the group leader, is achieved relatively easy. However, the ability to dynamically alter the distance which the group members are allowed to venture from the leader is not trivial; it hinges on the ability to correctly estimate distances within the wireless network which is something that I failed at. Thus I cannot provide practical proof that *all* of what is required for the solution proposed in this work is indeed also possible to realise. This does *not* mean that it is impossible to provide such a system. I am still very much of the opinion that an implementation of all the features I wanted to provide is achievable, given time. Ultimately, as my work on this thesis advanced, range and position estimation just proved to be beyond its scope, owing mostly to the higher than expected primary research to be done. My personal opinion, after having invested a lot of time into the prototype, however, remains that it is possible.

My work does facilitate some grounding theories as far as the proposed solution of an electronic assistance system for group management is concerned. Opinions on the idea that I have proposed were very favourable among the participants of the expert interviews, usability evaluations and field tests. Inquisitiveness for the project often went beyond polite asking "how things are coming along" and showed genuine interest. Both acceptance and demand for what I have proposed in this work were in ample supply. Having analysed the problematic situations and observed them first hand during the field trips, I conclude that an assistance system such as the one that I have proposed herein would be of great benefit to the working conditions of supervisory personnel as well as the safety of their charges.

Concerning the usability and usefulness of the devices: The usability and field tests have found the main device's functionality to be completely sufficient in order to carry out all tasks necessary for group management. I really see no need to bloat it beyond what is currently implemented of the user interface. Surely, the display could be made to be more appealing and the look of the menus and screens can be improved. Yet I would recommend against adding much functionality. What is there is sufficient: the possibility to add and remove members, to sound an alarm and display the name of missing members, a small set of options (mainly for sound and display lock), the possibility to send a signal to all group members and ideally the ability to alter the range on the fly. For interactions, even a small display and three buttons are sufficient. These are the features that I have found are required for the device to be useful, and any additions should be carefully weighed against the risk of making the interface too complicated to be quickly accessible. The point I am trying to make is that additional bells and whistles may do more harm than good in this context of usage. The proposed system will be used in highly stressful situations and should never make those worse by complicating interaction. For example, it is noted in [77, p. 3] during an examination of cognitive resource shortage in mobile human computer interactions how "*walking through a busy street to a bus stop is a taxing task requiring planning a route, managing time-to-target, and walking while at the same time taking care of safety*", even for a single person that is not tasked with the safekeeping of a whole group. Said work reveals how shortage of cognitive resources in a mobile context necessitates the user's attention on the device to be very fragmented and of a short-term nature; it is divided into four to eight second bursts and interrupted by the need to observe the environment or controlling mobility (e.g. walking).

For these good reasons, the prototype leaves only a minimal footprint on the user's cognitive resource pool. There is but one case where the device demands attention on its own accord: an alarm caused by stragglers requires acknowledgement in the form of a button-press in order to silence it. Most likely, the user will also take a look at the display to find out who is missing; this is somewhat logical and was observed during field trips. Since requiring a high level of visual attention is not viable in the mobile context [79, 58], all of the necessary information (who is missing) is always immediately visible. The only way to make the device take up less cognitive resources would be to move away from the idea of actively acknowledging the alarm, instead silencing it after a certain time period. This would facilitate the user to acknowledge it on his or her own choosing, but also harbours a danger: tasks "outside" of the device may demand attention more pressingly [58], making the user temporarily ignore it. This could result in the user effectively forgetting to check on (or plainly not hearing) the self-silencing notification in a situation where he or she has to cope with a lot of multitasking. I have witnessed, during the field trips, that the teacher finished other tasks most of the time before looking at the device to acknowledge the alarm. Those were short tasks, like telling someone to keep up or checking on a kid that lamented having gotten his clothes wet, yet they seemingly took priority over interacting with the device. Surely, this is true for other forms of mobile group management also.

I conclude, in accordance with my observations during field testing and the literature I have cited thus far, that the best behaviour for a device that aids in management of mobile groups is this: it should never demand attention unnecessarily or over long periods of time (which is, in this context, as little as four to eight seconds [77]), as group leaders usually have their hands full with other tasks. If it *does* need to draw the user's focus, it should do so in a way that lets him or her finish the current task before shifting attention, yet is nagging enough to require interaction sooner or later, and then display only the most important information, as pointedly as possible. For this context, I believe it is best to assume that the user is under high cognitive load – leave him or her alone unless it is absolutely necessary to get their attention, and if you do then say what you have to say as concise as possible.

I have made a point of preferring hardware buttons to a touch-based input on page 49. Those are also part of a user interface that aims to be less taxing on the user's attention. In 2000, Pascoe *et al.* used the term "Minimal Attention User Interface" and argue that "*shifting the human-computer interaction to unused channels or senses, and in a way that is not so cognitively demanding to distract the user from the task at hand*"[79, p 9] is a part thereof. They have shifted certain functions to hardware buttons during situations that require the user's attention elsewhere, stating tactile feedback and the possibility of eyes-free interaction as advantages. Using a touch based interface requires additional cognitive resources, because it necessitates that eye contact be maintained with the display and additional motoric function be coordinated in order to hit the right area [77]. By using hardware buttons, I have facilitated a form of interaction that is less demanding regarding the user's attention span. This design choice has worked well in field settings, and I strongly encourage future research to take into account the positive factors of physical buttons as opposed to the soft buttons found on touch-based input devices.

Something that is mentioned in both [79] and [58] is how it is preferred or necessitated in the mobile context that the user be able to interact with the device using a single hand only. The user may have his or her other hand occupied, which is not uncommon in field settings. The design of my prototype does facilitate this. Usability testing in a controlled environment has shown that, after a while, the participants started using both hands for interacting with the device (page 87). However, in the field this could not be observed. The device was, during both field trips, always carried and interacted with using only one hand; this is in accord with what was found in [79] and [58]. The teacher used the other hand for various activities like punching tickets when entering a tramway, ushering children forward or pointing things out. This is strong evidence that providing the possibility of one-handed usage is vital to usability in the field.

As far as the child devices are concerned, I have found their impact on the children's behaviour to be minimal. Observation has shown that their movement, play and demeanour during the field trips did not alter with or without wearing the items; this was confirmed by a teacher also. They were completely unconcerned with carrying foreign objects on their persons, an observation that is backed by the wear and tear that the devices exhibited after only two field trips. Practicality was found to be best with the high visibility vests. They may be used independently of weather conditions as the top layer of clothing and also provide good radio signal strength (which could be better if the transceiver was positioned more on the shoulder area than on the chest). The pendant also proved versatile as it may be carried around the neck without regard for other clothing. The only thing that makes it less preferable to the vests is the factor of visibility and that it is easier for a child to discard without the teacher noticing. The caps proved to be too reliant on just the right temperature to be comfortable; in hot weather they immediately fell from grace with all children. Furthermore their electronics proved to be the most finicky due to the limited space available. They are the only items I strongly discourage for future implementations.

It is notable that the children exhibited high acceptance of the items I have provided. When questioned as to whether or not they felt uneasy about carrying such devices as help their teacher localise them, not a single one of them seemed to find the notion unsettling. They seemed to accept without much afterthought that this was a good thing and helps their teacher, which in fact I found quite astounding considering their age (seven to eight years). Whether it should be alarming that they are so little concerned about the possibility of surveillance, if only locally, I leave up to the reader to decide.

Speaking of an aversion to surveillance, it is my personal opinion that the importance of using only local means for detecting absentees is invaluable to fostering acceptance towards such a system. I feel that this is a point that needs to be stressed: nothing has quite the potential to calm initial scepticism about this idea than dropping the fact that it only works in a limited radius around the teacher and does neither facilitate exact nor global localisation. As has been shown by the pilot project of Swedish daycare centres, public scepticism and doubt was very high when GPS was involved [11, 12], as was mistrust when individual localisation systems for families where used with a global range and exact resolution [32, 94]. Privacy concerns may largely be mitigated by a system such as was proposed and implemented in this work.

On the topic of field trips. On page 89 I have brought up the importance of field testing. Now, at this point in my work, after having gathered some experience both with testing in the field and in a controlled environment, I would like to elaborate on this. In the previously cited [72] and [43] field testing is regarded as very useful, as it helps identify problematic areas that were not obvious in a controlled environment. In [77] a point is also made for the importance of field tests, as the user's attention is much less fragmented under laboratory conditions. This backs the argument that field testing reveals flaws in human-computer interaction that may not be apparent in more "traditional" usability testing. Opposed to that is the stance taken in [56], where the authors argue that testing in the field did not bring forth much added value. Judging from the experience gathered during this work, I tend to agree with the former point of view. For me, field testing revealed some information and also confirmed some assumptions made during development. The evaluation of the child devices certainly took the most benefit out of it; I am not sure how their design and practicality (apart from the wireless functionality) could have been tested sufficiently otherwise. For the main device, field testing did not reveal as much additional information, yet there was some; the modification of the registration process was one of them, for example. It also confirmed the aptness of the interface, which is less of a new information, but rather an important confirmation.

I *do* however agree with one point made in [56], which is that the lack of control in field testing is certainly problematic. This is especially true when working with limited resources. Working alone, I did have severe difficulties processing and registering all of the information regarding usability and found it hard to maintain structure. Take, for example, the functionality of giving a signal to the children through the main device. This was never tested in the field because it did not fit the schedule during the first field test and on the second trip the child devices broke due to physical stress. I believe that the general problem of enacting control during field testing while not interfering too much with the observed can be eased to a certain amount, yet never completely, by experience. The second field test was already much easier for me in terms of information gathering and processing, simply because I knew what to expect.

Considering all of these factors, I deem field testing to be supplementary to a laboratory environment when one is working in the mobile context. During early iterations, usability testing under controlled conditions is the better choice without a doubt, but in the later stages of development it is, with a good probability, beneficial to take to the field at least once as it can reveal hitherto undiscovered aspects – I would not call field testing absolutely essential, but worth-wile with a good degree of likeliness.

Ultimately, what I have provided here for the reader is some fundamental work on the topic of assisting with the management of mobile groups through electronically aiding in awareness of presence. I have shown a prototype system that features a good set of functionalities, implemented after careful research of available technologies and requirement analysis. I have described how there are strong indications of demand for such a device and how the acceptance was very high among a group that is to be considered representative of prospective users. I have also given a report on what I think is best suited for a child device, based on observations and interviews done with a group of twenty five children over the course of two field tests. The implementation and evaluation of the prototype is to be viewed as an example use-case, a first exploration of this concept, and not as the one and only way of application. Surely, the proposed solution can be applied to a broader scope, such as tourist groups, skiing groups, *etc.* – essentially any other form of mobile group which requires supervision. If you take away nothing else from this work, then let it be that this topic seems a promising area for further research with a high acceptance and demand among the target audience.

5.2 Caveats

That all (meaning the conclusio) being said, there are some caveats to my work which I will not withhold. Firstly there is the matter of my technology assessment being done on paper only. I pieced together what I know of technologies presented in Section 2.2 from various sources and only compared by that. It is not that I do not trust these sources, but rather that they may have conducted their examinations in different environments, with different technological backgrounds and focuses, thus not providing an optimal basis for comparison. Ideally, I would have had to implement a prototype for each of the likely technologies (WiFi, Bluetooth, ZigBee and if at all possible something utilising Ultra Wideband) and compared them to each other in identical environments with the same premises. As it stands, this is beyond the scope of my work as far as both time and money are concerned, and thus my choice of technology was based on possibly slanted comparisons. Maybe WiFi would indeed have provided better connection features? Only a more in-depth comparison in an actual implementation scenario would show.

Furthermore there is the fact that I did only implement and test a prototype with five group members. An actual system as would adhere to the problem solution discussed in this work would have had to have at least 25 child devices; one for each child of the class. With so small a number of devices in the network, it may be that certain problems with the setup did not become obvious, like for example those concerning network traffic load. I cannot say if the main device, in the state that it is currently in, could have handled simultaneous updates from so many group members. I have also stated that mesh networking should overcome the high interferences found in many applications scenarios; I assume thus, but I have no way to be sure without a full-scale implementation. Hence there were some presumptions involved when I stated that I believe it possible to implement my proposed solution.

Lastly, one may call into question the fidelity of the prototype and whether or not it was really good enough to be tested in the field. As I have said in the previous section, I deemed the prototype to be of high enough fidelity to be put on trial in a real working environment. Calling it high fidelity would probably be pushing it a bit, but high *enough* is not that far-fetched I think. However, as there exist various perceptions of the notion of prototype fidelity, opinions on that may differ. The question of what constitutes a low and high fidelity prototype, what advantages these approaches hold and which one is best suited (to the mobile environment) is a sprawling topic with many converse opinions and theses floating around, which I will not go into at length. Personally, I think the binary categorisation into high and low is insufficient; rather, I find the multi dimensional approach as discussed in e.g. [96] and [69] to be most accommodating. Both works argue that a prototype can be of high fidelity in some areas and low fidelity in others; they diverge somewhat on the definition of what those dimensions are. They have in common to distinguish between the number of features available, the degree of functionality which they offer, the presence of visual aesthetics in the prototype as well as the quality of interactions. I consider the prototype to be of high fidelity in the compartments feature functionality as well as quality of interactions, yet lacking in overall device aesthetics and polish of the interface as well as providing medium feature breadth (as it is currently lacking range and position estimations). According to [96], this alone is enough to classify it as low fidelity whereas [69] speaks of "mixed-fidelity" in favour of dropping the traditional, binary categorisation altogether. I consider my prototype to be "sufficient", as it provides much of the features to a high degree of functionality and offers a well developed interaction scheme. And yet it is to be regarded as low fidelity by certain standards, and thus may well be considered unfit for field testing by some.

5.3 Future Work

As was mentioned in Section 2.3, the approach described in [105] shows promise for this field of application. It estimates node positions in a relative coordinate system, which is all that is really needed for my project. It would be a great help for the teacher if he or she would know the position of the children relative to that of the main device. This would indicate "where to look" for members of the group. The authors claim movement speeds under 20 metres per second are supported, so mobility should not be a problem. Rather, the challenging part would be to obtain distance measurements that are accurate enough even in an interference-rich environment; in [105], *Time of Arrival* based calculations were used. These were not possible to implement with the prototyping platform of my choice as was explained in Section 3.3. A goal for future work should be to do a more in-depth investigation of the possibilities of range and location estimation for a completely mobile wireless sensor network, without anchors. With many other areas of interest to be covered in this work, I found those to be beyond its scope eventually1. For example, the ZigBee standards (and, as an implementation thereof, also the XBee modules) provide something called *ZigBee Device Objects* (ZDO) on the application layer, which offer significant network management capabilities. Those include reading any device's neighbouring table which, besides network addresses, also holds LQI readings; that offers a good option to gauge the network's topology. Certainly, some more time to experiment with such possibilities could yield promising results. A working location system or even range estimates would greatly improve any implementation. That I could not provide these is more than just a little vexing.

Another topic that should be investigated is the long term acceptance and handling of the devices. I only accompanied school classes on two field trips. During these, the children were reasonably positive towards the items I gave them and wore those quite happily. The questions that remain are whether this will still be the case after they had to wear them on every school outing for a year or two; or four, for that matter. Would they get tired of wearing these items and try to discard them without the teacher's consent? Would it change the way they act and behave during field trips? How about the teachers, how would their behaviour change after truly using such an electronic aid for a few years? Would they stop counting for absentees manually, or do they harbour distrust in technological assistance? Would their acceptance and positive attitude remain? These are questions that require a long term study which is beyond the scope of this work, but should be of interest when further investigating the impacts of such electronic assistance systems. A topic which does not require a long term study, yet was omitted in this work, is that of the parent's acceptance towards my proposed solution. As I have said in Section 4.2, all but very few agreed to their child partaking in the field test, yet this is not indicative of the reception of a real system being put into place. I expect a higher acceptance than with global positioning solutions, such as were used in Swedish daycare centres or for tracking family members, once the fact that presence detection works locally only is made abundantly clear to the affected. But this is only speculation – the topic needs to be investigated in a proper manner.

Furthermore I would wish for a full scale prototype implementation, including upwards of 25 child devices and maybe a second main device. The network's behaviour is bound to change when a larger number of devices are used, bringing into the game such notions as load (both of the network and the processing power of the main devices) and mesh networking dynamics that did not really occur with the testing of only five items. I have shown that it is possible to implement my solution on a small scale, strongly inferring that larger groups are also well within the possible. But I did not bring forth irrefutable proof.

In Section 5.1 I have mentioned eyes-free interaction. This may also be of interest for further iterations. I have deliberately facilitated it (or rather, the possibility thereof) through the use of hardware buttons yet admittedly the only task that makes use of this is acknowledging an alarm. As interaction without visual contact to the screen would mean less usage of cognitive resources, I would think it highly desirable. An example may be that the group leader could send a signal to all member by pressing the central button (or any other) for a longer time, significantly speeding up the interaction and doing away with the need to look at the display. I am sure that other such improvements could be made to facilitate eyes-free interactions.

As I have stated on page 73 and following, I do not currently see the need to implement extensive security measures for this system. However, if somebody were to take up this work and wanted to focus more heavily on that topic, the possibilities are manifold. Since security is just about non-existent in my prototype, everything remains to be done still in this regard. Improvements could range from encrypting communications to establishing a secure system for distributing private keys to implementing some form of packet spoofing detection to pretty much everything you could read in any book about security in wireless networks. None of that has been done, and remains open for future investigations if so desired.

5.4 Acknowledgement

I want to give thanks to my mother and grandmothers who have helped me a lot during the implementation of the child devices and certainly deserve acknowledgement. Their advice and assistance proved invaluable during the tailoring of various textiles and clothing. Without them, it would not have been possible to provide such high quality items for housing the electronics. Similarly, I have to thank my grandfather, for enabling me to bring forth a custom-made chassis for the main device. He has helped me tremendously during the whole process of adapting the casing and fitting the electronics inside. Without his aid the main device would have looked a lot less high-fidelity than it does. Thanks to these people, the prototypes look like they do – I would never have managed to accomplish such work without them. Thanks goes also to my significant other's mother, for taking the time to proof read this whole wall of text.

I feel that acknowledgement is also due for my supervisor who managed to make me get stuff done. As I heard, it is not to be taken for granted nowadays to have a supervisor that gives actual, honest feedback and takes an active interest in their student's work, such as had I.

Lastly, I want to thank my significant other herself, for bearing with me while I was doing this project, even as it dragged on: you always accepted how absorbed I was in my work, without holding a grudge. I love you, very much.

Bibliography

[1] 2011 revision of the ieee 802.15.4 standard. http://standards.ieee.org/getieee802/download/802.15.4-2011.pdf. Visited on August 11th, 2013.

[2] 2011 revision of the ieee 802.15.4 standard. http://www.digi.com/products/wireless-wired-embedded-solutions/zigbee-rf-modules/zigbee-mesh-module/xbee-zb-module#docs. Visited on August 11th, 2013.

[3] Arduino library for adressing the sparkfun lcd shield. https://github.com/TCWORLD/gLCD-Library. Visited on July 30th, 2013.

[4] Arduino library for communicating with xbees. https://code.google.com/p/xbee-arduino/. Visited on July 30th, 2013.

[5] Bluetooth special interest group annual report 2011 overview. https://www.bluetooth.org/About/AnnualReport/AnnualReport11/Overview.htm. Visited on September 8th, 2012.

[6] Galileo satellite launches. http://ec.europa.eu/enterprise/policies/satnav/galileo/satellite-launches/index_en.htm. Visited on September 20th, 2012.

[7] Ieee 802.11 homepage. http://www.ieee802.org/11/. Visited on October 13th, 2012.

[8] lok8u homepage. http://www.lok8u.com/. Visited on May 3rd, 2012.

[9] "mobile terminal group management system and mobile terminal" patent. http://worldwide.espacenet.com/publicationDetails/biblio;jsessionid=B987387B8C60A4F79012FE9471EADF8B.espacenet_levelx_prod_4?FT=D&date=20111222&DB=worldwide.espacenet.com&locale=en_EP&CC=JP&NR=2011259056A&KC=A&ND=5. Visited on May 16th, 2012.

[10] Official arduino project homepage. http://www.arduino.cc/. Visited on July 28th, 2013.

[11] online article about a tracking system used in daycare centres in sweden. http://www.stuff.co.nz/technology/digital-living/5669941/Daycare-centres-use-GPS-to-track-children. Visited on August 15th, 2013.

[12] online article about a tracking system used in daycare centres in sweden. http://www.google.com/hostednews/afp/article/ALeqM5hiCFhm_xOq6TKOiG6hHbc7kqn0TQ. Visited on August 15th, 2013.

[13] Pnt information homepage. http://www.gps.gov/governance/excom/nco/. Visited on September 25th, 2012.

[14] Pulson 410 rcm. http://www.timedomain.com. Visited on October 14th, 2012.

[15] Softwareserial library for serial communication on arbitrary arduino pins. http://arduino.cc/en/Reference/SoftwareSerial. Visited on August 3rd, 2013.

[16] Sparkfun arduino xbee 2 module. https://www.sparkfun.com/products/11217. Visited on October 15th, 2012.

[17] Standard c++ "port" for arduino. https://github.com/maniacbug/StandardCplusplus. Visited on July 30th, 2013.

[18] Toddler tag child locator. http://www.brickhousesecurity.com/child-locator.html. Visited on May 3rd, 2012.

[19] ubisafe homepage. http://www.UbiSafe.com/. Visited on May 3rd, 2012.

[20] Wifi alliance article. http://www.wi-fi.org/knowledge-center/articles/did-you-know. Visited on October 13th, 2012.

[21] Zebra uwb dart tag (not actually for tagging zebras). http://www.zebra.com. Visited on October 14th, 2012.

[22] Zigbee homepage – about. http://zigbee.org/About/AboutTechnology/ZigBeeTechnology.aspx. Visited on October 13th, 2012.

[23] Zigbee homepage – shipments. http://zigbee.org/Default.aspx?Contenttype=ArticleDet&tabID=332&moduleId=806&Aid=410&PR=PR1. Visited on October 13th, 2012.

[24] Zigbee homepage – standards. http://zigbee.org/About/AboutTechnology/Standards.aspx. Visited on October 13th, 2012.

[25] Official arduino project homepage. `http://spectrum.ieee.org/geek-life/hands-on/the-making-of-arduino/`, 2011. Visited on July 28th, 2013.

[26] J. Albowicz, A. Chen, and Lixia Zhang. Recursive position estimation in sensor networks. In *Network Protocols, 2001. Ninth International Conference on*, pages 35–41, 2001.

[27] N. Baker. Zigbee and bluetooth strengths and weaknesses for industrial applications. *Computing Control Engineering Journal*, 16(2):20 –25, april-may 2005.

[28] Paolo Baronti, Prashant Pillai, Vince W.C. Chook, Stefano Chessa, Alberto Gotta, and Y. Fun Hu. Wireless sensor networks: A survey on the state of the art and the 802.15.4 and zigbee standards. *Computer Communications*, 30(7):1655 – 1695, 2007.

[29] Fehmi Ben Abdesslem, Andrew Phillips, and Tristan Henderson. Less is more: energy-efficient mobile sensing with senseless. In *Proceedings of the 1st ACM workshop on Networking, systems, and applications for mobile handhelds*, MobiHeld '09, pages 61–62, New York, NY, USA, 2009. ACM.

[30] N.W. Bergmann, M. Wallace, and E. Calia. Low cost prototyping system for sensor networks. In *Intelligent Sensors, Sensor Networks and Information Processing (ISSNIP), 2010 Sixth International Conference on*, pages 19–24, 2010.

[31] J. Blumenthal, R. Grossmann, F. Golatowski, and D. Timmermann. Weighted centroid localization in zigbee-based sensor networks. In *Intelligent Signal Processing, 2007. WISP 2007. IEEE International Symposium on*, pages 1 –6, oct. 2007.

[32] Julie Boesen, Jennifer A. Rode, and Clara Mancini. The domestic panopticon: location tracking in families. In *Proceedings of the 12th ACM international conference on Ubiquitous computing*, Ubicomp '10, pages 65–74, New York, NY, USA, 2010. ACM.

[33] A. Boukerche, H.A.B. Oliveira, E.F. Nakamura, and A.A.F. Loureiro. Localization systems for wireless sensor networks. *Wireless Communications, IEEE*, 14(6):6–12, 2007.

[34] Andrew Bragdon, Eugene Nelson, Yang Li, and Ken Hinckley. Experimental analysis of touch-screen gesture designs in mobile environments. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, CHI '11, pages 403–412, New York, NY, USA, 2011. ACM.

[35] Jennifer Bray. *Bluetooth 1.1: connect without cables*. Prentice-Hall, Inc., 2nd edition, 2002.

[36] J. Cache, J. Wright, and V. Liu. *Hacking Exposed Wireless, Second Edition*. Hacking Exposed. Mcgraw-hill, 2010.

[37] Aaron Carroll and Gernot Heiser. An analysis of power consumption in a smartphone. In *Proceedings of the 2010 USENIX conference on USENIX annual technical conference*, USENIXATC'10, pages 21–21, Berkeley, CA, USA, 2010. USENIX Association.

[38] G. Chandrasekaran, M.A. Ergin, Jie Yang, Song Liu, Yingying Chen, M. Gruteser, and R.P. Martin. Empirical evaluation of the limits on localization using signal strength. In *Sensor, Mesh and Ad Hoc Communications and Networks, 2009. SECON '09. 6th Annual IEEE Communications Society Conference on*, pages 1 –9, june 2009.

[39] M.V.S. Chandrashekhar, P. Choi, K. Maver, R. Sieber, and K. Pahlavan. Evaluation of interference between ieee 802.11b and bluetooth in a typical office environment. In *Personal, Indoor and Mobile Radio Communications, 2001 12th IEEE International Symposium on*, volume 1, pages D–71 –D–75 vol.1, sep 2001.

[40] Chen-YQ. Chen Wu Ding X.L. Li Z.L. Wang N. Chao, J. and M. Yu. An experimental investigation into the performance of gps based vehicle positioning in very dense urban areas. *Journal of Geospatial Engineering*, 3(1):59–66, 2001.

[41] Irena Pletikosa Cvijikj and Florian Michahelles. The toolkit approach for end-user participation in the internet of things. In Dieter Uckelmann, Mark Harrison, and Florian Michahelles, editors, *Architecting the Internet of Things*, pages 65–96. Springer Berlin Heidelberg, 2011.

[42] Hans Dodel and Dieter Häupler. *Satellitennavigation*. Springer Verlag Berlin Heidelberg, 2nd revised and extended edition, 2010.

[43] Henry Been-Lirn Duh, Gerald C. B. Tan, and Vivian Hsueh-hua Chen. Usability evaluation for mobile device: a comparison of laboratory and field tests. In *Proceedings of the 8th conference on Human-computer interaction with mobile devices and services*, MobileHCI '06, pages 181–186, New York, NY, USA, 2006. ACM.

[44] Shahin Farahani. *Zigbee wireless networks and transceivers*. Elsevier Ltd., 1st edition, 2008.

[45] R.J. Fontana. Recent system applications of short-pulse ultra-wideband (uwb) technology. *Microwave Theory and Techniques, IEEE Transactions on*, 52(9):2087 – 2104, sept. 2004.

[46] Matthew S. Gast. *802.11 Wireless Networks: The Definitive Guide*. O'Reilly Media, Inc., 2nd edition, 2005.

[47] K. Gill, Shuang-Hua Yang, Fang Yao, and Xin Lu. A zigbee-based home automation system. *Consumer Electronics, IEEE Transactions on*, 55(2):422 –430, may 2009.

[48] Ramakrishna Gummadi, David Wetherall, Ben Greenstein, and Srinivasan Seshan. Understanding and mitigating the impact of rf interference on 802.11 networks. *SIGCOMM Comput. Commun. Rev.*, 37(4):385–396, August 2007.

[49] Wenqi Guo, W.M. Healy, and Mengchu Zhou. Zigbee-wireless mesh networks for building automation and control. In *Networking, Sensing and Control (ICNSC), 2010 International Conference on*, pages 731 –736, april 2010.

[50] Wenqi Guo, W.M. Healy, and MengChu Zhou. Impacts of 2.4-ghz ism band interference on ieee 802.15.4 wireless sensor network reliability in buildings. *Instrumentation and Measurement, IEEE Transactions on*, 61(9):2533 –2544, sept. 2012.

[51] Mathias Hein and Dr. Bernd Maciejewski. *Wireless LAN - Funknetze in der Praxis*. Franzis Verlag GmbH Poing, 1st edition, 2003.

[52] Alan Holt and Chi-Yu Huang. *802.11 Wireless Networks*. Springer London, 1st edition, 2010.

[53] Gerhard Kafka. *WLAN - Technick, Standards, Planung und Sicherheit für Wireless LAN*. Carl Hanser Verlag München Wien, 1st edition, 2005.

[54] A. Kamerman and N. Erkocevic. Microwave oven interference on wireless lans operating in the 2.4 ghz ism band. In *Personal, Indoor and Mobile Radio Communications, 1997. 'Waves of the Year 2000'. PIMRC '97., The 8th IEEE International Symposium on*, volume 3, pages 1221 –1227 vol.3, sep 1997.

[55] Elliot Kaplan and Christopher Hegarty. *Understanding GPS: principles and applications*. Artech House, Inc., 2nd edition, 2006.

[56] Jesper Kjeldskov, MikaelB. Skov, BenedikteS. Als, and RuneT. Høegh. Is it worth the hassle? exploring the added value of evaluating the usability of context-aware mobile systems in the field. In Stephen Brewster and Mark Dunlop, editors, *Mobile Human-Computer Interaction - MobileHCI 2004*, volume 3160 of *Lecture Notes in Computer Science*, pages 61–73. Springer Berlin Heidelberg, 2004.

[57] A. Koubaa, M. Ben Jamaa, and A. AlHaqbani. An empirical analysis of the impact of rss to distance mapping on localization in wsns. In *Communications and Networking (ComNet), 2012 Third International Conference on*, pages 1 –7, 29 2012-april 1 2012.

[58] Steinar Kristoffersen and Fredrik Ljungberg. Making place to make it work: empirical explorations of hci for mobile cscw. In *Proceedings of the international ACM SIGGROUP conference on Supporting group work*, GROUP '99, pages 276–285, New York, NY, USA, 1999. ACM.

[59] Gerald Kupris and Axel Sikora. *Zigbee: Datenfunk mit IEEE 802.15.4 und ZigBee*. Franzis Verlag GmbH, 1st edition, 2007.

[60] Kaushik Lakshminarayanan, Srinivasan Seshan, and Peter Steenkiste. Understanding 802.11 performance in heterogeneous environments. In *Proceedings of the 2nd ACM SIGCOMM workshop on Home networks*, HomeNets '11, pages 43–48, New York, NY, USA, 2011. ACM.

[61] Seng-Yong Lau, Tsung-Han Lin, Te-Yuan Huang, I-Hei Ng, and Polly Huang. A measurement study of zigbee-based indoor localization systems under rf interference. In *Proceedings of the 4th ACM international workshop on Experimental evaluation and characterization*, WINTECH '09, pages 35–42, New York, NY, USA, 2009. ACM.

[62] Jin-Shyan Lee, Yu-Wei Su, and Chung-Chou Shen. A comparative study of wireless protocols: Bluetooth, uwb, zigbee, and wi-fi. In *Industrial Electronics Society, 2007. IECON 2007. 33rd Annual Conference of the IEEE*, pages 46 –51, nov. 2007.

[63] Hongwei Li, Zhongning Jia, and Xiaofeng Xue. Application and analysis of zigbee security services specification. In *Networks Security Wireless Communications and Trusted Computing (NSWCTC), 2010 Second International Conference on*, volume 2, pages 494–497, 2010.

[64] Nia-Chiang Liang, Ping-Chieh Chen, Tony Sun, Guang Yang, Ling-Jyh Chen, and M. Gerla. Impact of node heterogeneity in zigbee mesh network routing. In *Systems, Man and Cybernetics, 2006. SMC '06. IEEE International Conference on*, volume 1, pages 187–191, 2006.

[65] Siriwongpairat-W. Liu, K. *Ultra-Wideband Communications Systems :Multiband OFDM Approach*. Wiley-IEEE Press, 2008.

[66] Y. Luo and C. Law. Indoor positioning using uwb-ir signals in the presence of dense multipath with path overlapping. *Wireless Communications, IEEE Transactions on*, PP(99):1 –10, 2012.

[67] H. Mazar. A comparison between european and north american wireless regulations. In *Telecom World (ITU WT), 2011 Technical Symposium at ITU*, pages 182 –186, oct. 2011.

[68] D.D. McCrady, L. Doyle, H. Forstrom, T. Dempsey, and M. Martorana. Mobile ranging using low-accuracy clocks. *Microwave Theory and Techniques, IEEE Transactions on*, 48(6):951–958, 2000.

[69] Michael McCurdy, Christopher Connors, Guy Pyrzak, Bob Kanefsky, and Alonso Vera. Breaking the fidelity barrier: an examination of our current characterization of prototypes and an example of a mixed-fidelity success. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, CHI '06, pages 1233–1242, New York, NY, USA, 2006. ACM.

[70] E. Mok, G. Retscher, and L. Lau. Development of an event-reporting intelligent vehicle navigation system for areas with urban canyons. *Survey Review*, 36(284):398–409, 2002.

[71] Tom Nicolai and Holger Kenn. About the relationship between people and discoverable bluetooth devices in urban environments. In *Proceedings of the 4th international conference on mobile technology, applications, and systems and the 1st international symposium on Computer human interaction in mobile technology*, Mobility '07, pages 72–78, New York, NY, USA, 2007. ACM.

[72] Christian Monrad Nielsen, Michael Overgaard, Michael Bach Pedersen, Jan Stage, and Sigge Stenild. It's worth the hassle!: the added value of evaluating the usability of mobile systems in the field. In *Proceedings of the 4th Nordic conference on Human-computer interaction: changing roles*, NordiCHI '06, pages 272–280, New York, NY, USA, 2006. ACM.

[73] J. Nielsen. *Usability Engineering*. Interactive Technologies. Elsevier Science, 1993.

[74] Homayoun Nikookar and Ramjee Prasad. *Introduction to Ultra Wideband for Wireless Communications*. Springer Netherlands, 2009.

[75] S. Ondrej, B. Zdenek, F. Petr, and H. Ondrej. Zigbee technology and device design. In *Networking, International Conference on Systems and International Conference on Mobile Communications and Learning Technologies, 2006. ICN/ICONS/MCL 2006. International Conference on*, page 129, april 2006.

[76] I. Oppermann, L. Stoica, A. Rabbachin, Z. Shelby, and J. Haapola. Uwb wireless sensor networks: Uwen - a practical example. *Communications Magazine, IEEE*, 42(12):S27 – S32, dec. 2004.

[77] Antti Oulasvirta, Sakari Tamminen, Virpi Roto, and Jaana Kuorelahti. Interaction in 4-second bursts: the fragmented nature of attentional resources in mobile hci. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, CHI '05, pages 919–928, New York, NY, USA, 2005. ACM.

[78] Qixiang Pang and V.C.M. Leung. Channel clustering and probabilistic channel visiting techniques for wlan interference mitigation in bluetooth devices. *Electromagnetic Compatibility, IEEE Transactions on*, 49(4):914 –923, nov. 2007.

[79] Jason Pascoe, Nick Ryan, and David Morse. Using while moving: Hci issues in fieldwork environments. *ACM Trans. Comput.-Hum. Interact.*, 7(3):417–437, September 2000.

[80] N. Patwari, J.N. Ash, S. Kyperountas, A.O. Hero, R.L. Moses, and N.S. Correal. Locating the nodes: cooperative localization in wireless sensor networks. *Signal Processing Magazine, IEEE*, 22(4):54–69, 2005.

[81] Graham C. Kjeldskov J. Pedell, S. and J. Davies. Mobile evaluation: What the data and the metadata told us. In *Proceedings of OzCHI 2003*, volume 8, pages 96 –105, 2003.

[82] D. Porcino and W. Hirt. Ultra-wideband radio technology: potential and challenges ahead. *Communications Magazine, IEEE*, 41(7):66 – 74, july 2003.

[83] Heitor S. Ramos, Tao Zhang, Jie Liu, Nissanka B. Priyantha, and Aman Kansal. Leap: a low energy assisted gps for trajectory-based services. In *Proceedings of the 13th international conference on Ubiquitous computing*, UbiComp '11, pages 335–344, New York, NY, USA, 2011. ACM.

[84] Günther Retscher and Michael Thienelt. NAVIO – A Navigation and Guidance Service for Pedestrians. *Journal of Global Positioning Systems*, 3(1–2):208–217, 2004.

[85] S. Safaric and K. Malaric. Zigbee wireless standard. In *Multimedia Signal Processing and Communications, 48th International Symposium ELMAR-2006 focused on*, pages 259 –262, june 2006.

[86] Nel Samama. *Global Positioning: Technologies and Performance*. John Wiley & Sons, Inc., Hoboken, New Jersey, 1st edition, 2008.

[87] Martin Sauter. *Grundkurs Mobile Kommunikationssysteme*. Friedr. Vieweg & Sohn Verlag — GWV Fachverlage GmbH, Wiesbaden, 3rd, extended edition, 2008.

[88] Andreas Savvides, Chih-Chieh Han, and Mani B. Strivastava. Dynamic fine-grained localization in ad-hoc networks of sensors. In *Proceedings of the 7th annual international conference on Mobile computing and networking*, MobiCom '01, pages 166–179, New York, NY, USA, 2001. ACM.

[89] R.A. Scholtz, R. Weaver, E. Homier, J. Lee, P. Hilmes, A. Taha, and R. Wilson. Uwb radio deployment challenges. In *Personal, Indoor and Mobile Radio Communications, 2000. PIMRC 2000. The 11th IEEE International Symposium on*, volume 1, pages 620 –625 vol.1, 2000.

[90] Soo Young Shin, Hong Seong Park, Sunghyun Choi, and Wook Hyun Kwon. Packet error rate analysis of zigbee under wlan and bluetooth interferences. *Wireless Communications, IEEE Transactions on*, 6(8):2825 –2830, august 2007.

[91] K. Shuaib, M. Boulmalf, F. Sallabi, and A. Lakas. Co-existence of zigbee and wlan, a performance study. In *Wireless Telecommunications Symposium, 2006. WTS '06*, pages 1 –6, april 2006.

[92] M.L. Sichitiu and V. Ramadurai. Localization of wireless sensor networks with a mobile beacon. In *Mobile Ad-hoc and Sensor Systems, 2004 IEEE International Conference on*, pages 174–183, 2004.

[93] A. Sikora and V.F. Groza. Coexistence of ieee802.15.4 with other systems in the 2.4 ghz-ism-band. In *Instrumentation and Measurement Technology Conference, 2005. IMTC 2005. Proceedings of the IEEE*, volume 3, pages 1786 –1791, may 2005.

[94] Asimina Vasalou, Anne-Marie Oostveen, and Adam N. Joinson. A case study of non-adoption: the values of location tracking in the family. In *Proceedings of the ACM 2012 conference on Computer Supported Cooperative Work*, CSCW '12, pages 779–788, New York, NY, USA, 2012. ACM.

[95] Niko Vidgren, Keijo Haataja, Jose Luis Patino-Andres, Juan Jose Ramirez-Sanchis, and Pekka Toivanen. Security threats in zigbee-enabled systems: Vulnerability evaluation, practical experiments, countermeasures, and lessons learned. In *System Sciences (HICSS), 2013 46th Hawaii International Conference on*, pages 5132–5138, 2013.

[96] Robert A. Virzi, Jeffrey L. Sokolov, and Demetrios Karis. Usability problem identification using both low- and high-fidelity prototypes. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, CHI '96, pages 236–243, New York, NY, USA, 1996. ACM.

[97] M.Z. Win and R.A. Scholtz. On the robustness of ultra-wide bandwidth signals in dense multipath environments. *Communications Letters, IEEE*, 2(2):51 –53, feb. 1998.

[98] Jingjing Xia, Choi Look Law, Yuan Zhou, and Kee Siang Koh. 3 - 5 ghz uwb impulse radio transmitter and receiver mmic optimized for long range precision wireless sensor networks. *Microwave Theory and Techniques, IEEE Transactions on*, 58(12):4040 –4051, dec. 2010.

[99] Guochang Xu. *GPS - Theory, Algorithms and Applications*. Springer-Verlag Berlin Heidelberg New York, 1st edition, 2003.

[100] Liuqing Yang and G.B. Giannakis. Ultra-wideband communications: an idea whose time has come. *Signal Processing Magazine, IEEE*, 21(6):26 – 54, nov. 2004.

[101] Li Zhao and A.M. Haimovich. Performance of ultra-wideband communications in the presence of interference. *Selected Areas in Communications, IEEE Journal on*, 20(9):1684 – 1691, dec 2002.

[102] Li Zhao, A.M. Haimovich, and H. Grebel. Performance of ultra-wideband communications in the presence of interference. In *Communications, 2001. ICC 2001. IEEE International Conference on*, volume 10, pages 2948 –2952 vol.10, 2001.

[103] Y. Zhao, B.G. Agee, and J.H. Reed. Simulation and measurement of microwave oven leakage for 802.11 wlan interference management. In *Microwave, Antenna, Propagation and EMC Technologies for Wireless Communications, 2005. MAPE 2005. IEEE International Symposium on*, volume 2, pages 1580 –1583 Vol. 2, aug. 2005.

[104] Gang Zhou, Tian He, Sudha Krishnamurthy, and John A. Stankovic. Impact of radio irregularity on wireless sensor networks. In *Proceedings of the 2nd international conference on Mobile systems, applications, and services*, MobiSys '04, pages 125–138, New York, NY, USA, 2004. ACM.

[105] Srdjan Čapkun, Maher Hamdi, and Jean-Pierre Hubaux. Gps-free positioning in mobile ad hoc networks. *Cluster Computing*, 5(2):157–167, 2002.