Capacitive Sensing and Communication for Ubiquitous Interaction and Environmental Perception

dem Fachbereich Informatik
der Technischen Universität Darmstadt
genehmigte

DISSERTATION

zur Erlangung des akademischen Grades eines
Doktor-Ingenieurs (Dr.-Ing.)
von

M. Sc. Tobias Alexander Große-Puppendahl
geboren in Münster (Westfalen), Deutschland

Referenten der Arbeit:	Prof. Dr. techn. Dieter W. Fellner Technische Universität Darmstadt Prof. Dr. Kristof van Laerhoven Universität Freiburg
Tag der Einreichung:	24.03.2015
Tag der mündlichen Prüfung:	22.05.2015

Darmstädter Dissertation 2015
D 17

Erklärung zur Dissertation

Hiermit versichere ich die vorliegende Dissertation selbständig nur mit den angegebenen Quellen und Hilfsmitteln angefertigt zu haben. Alle Stellen, die aus Quellen entnommen wurden, sind als solche kenntlich gemacht. Diese Arbeit hat in gleicher oder ähnlicher Form noch keiner Prüfungsbehörde vorgelegen.

Darmstadt, den 24.03.2015 Tobias Alexander Große-Puppendahl

Abstract

During the last decade, the functionalities of electronic devices within a living environment constantly increased. Besides the personal computer, now tablet PCs, smart household appliances, and smartwatches enriched the technology landscape. The trend towards an ever-growing number of computing systems has resulted in many highly heterogeneous human-machine interfaces. Users are forced to adapt to technology instead of having the technology adapt to them. Gathering context information about the user is a key factor for improving the interaction experience. Emerging wearable devices show the benefits of sophisticated sensors which make interaction more efficient, natural, and enjoyable. However, many technologies still lack of these desirable properties, motivating me to work towards new ways of sensing a user's actions and thus enriching the context. In my dissertation I follow a human-centric approach which ranges from sensing hand movements to recognizing whole-body interactions with objects.

This goal can be approached with a vast variety of novel and existing sensing approaches. I focused on perceiving the environment with quasi-electrostatic fields by making use of capacitive coupling between devices and objects. Following this approach, it is possible to implement interfaces that are able to recognize gestures, body movements and manipulations of the environment at typical distances up to 50 cm. These sensors usually have a limited resolution and can be sensitive to other conductive objects or electrical devices that affect electric fields. The technique allows for designing very energy-efficient and high-speed sensors that can be deployed unobtrusively underneath any kind of non-conductive surface. Compared to other sensing techniques, exploiting capacitive coupling also has a low impact on a user's perceived privacy.

In this work, I also aim at enhancing the interaction experience with new perceptional capabilities based on capacitive coupling. I follow a bottom-up methodology and begin by presenting two low-level approaches for environmental perception. In order to perceive a user in detail, I present a rapid prototyping toolkit for capacitive proximity sensing. The prototyping toolkit shows significant advancements in terms of temporal and spatial resolution. Due to some limitations, namely the inability to determine the identity and fine-grained manipulations of objects, I contribute a generic method for communications based on capacitive coupling. The method allows for designing highly interactive systems that can exchange information through air and the human body. I furthermore show how human body parts can be recognized from capacitive proximity sensors. The method is able to extract multiple object parameters and track body parts in real-time. I conclude my thesis with contributions in the domain of context-aware devices and explicit gesture-recognition systems.

Zusammenfassung

Innerhalb der letzten Jahre haben sich die Funktionalitäten der Geräte in einer Wohnumgebung stark erweitert. Dieser Trend kann zu großen Teilen dem technischen Fortschritt in der Elektro- und Informationstechnik zugeschrieben werden. Im Laufe dieser Entwicklung wurden bestehende Geräte, wie bspw. Personal Computer, miniaturisiert, und neue Nutzungsszenarien entstanden. Durch die Verfügbarkeit von immer effizienteren Prozessoren und neuen energiesparenden Sensoren entwickelten sich neue Gerätekatgorien wie Smartwatches und Smartphones. Neben diesen vielen unterschiedlichen Geräten existieren ebenfalls viele heterogene Benutzungskonzepte. Die Nutzung der Geräte ist häufig mit Problemen innerhalb aller Altersklassen verbunden, da immer mehr Nutzungsarten verinnerlicht werden müssen. Obwohl die technischen Möglichkeiten moderner Geräte meist sehr groß sind, sind Nutzer oft nur in der Lage einen kleinen Anteil tatsächlich auszuschöpfen.

Aufgrund dieser Probleme entstanden neue Interaktionsarten, bspw. basierend auf Gesten oder Sprache. Die damit verbundenen neuen Technologien haben ein gemeinsames Ziel - sie sollen Interaktion natürlicher und einfacher machen. Aber auch solche Interaktionsarten führen zu neuem Lernaufwand bei Nutzern und somit zunächst zu erhöhter Komplexität. Trotzdem können diese natürlicheren Formen der Mensch-Technik-Interaktion als ein Schritt in die richtige Richtung angesehen werden.

Der Weg zu besser generalisierbarem und verständlichen Interaktionsdesign ist eine große Herausforderung für die Wissenschaft. Neben einer höheren Interoperabilität zwischen unterschiedlichen Geräten kann ein Verständnis des Nutzers und seiner aktuellen Situation die Interaktion mit einem technischen System wesentlich vereinfachen [Sch00]. Nach einem bekannten Artikel von Mark Weiser [Wei99] kann eine intelligente Umgebung, die die Nutzerwünsche erkennt und die Erfüllung dieser unterstützt, als Ziel angesehen werden. In einer solchen Welt integrieren sich Geräte unauffällig und eigenständig in das technische System. Somit verschwindet die Technologie von der Blickfläche und wird ein nicht unterscheidbarer Teil der Umgebung eines Nutzers. 15 Jahre nachdem diese Vision formuliert wurde, hat sie nichts von ihrer Relevanz verloren. Auch wenn einige Geräteklassen, wie Smartwatches, bereits sehr intelligent sind, ist die Entwicklung noch nicht in der breiten Masse der eingebetteten Systeme in einer Wohnumgebung angekommen.

Einer der wichtigsten Faktoren, um Weisers Ziel zu erreichen, ist die Wahrnehmung des Nutzers und seiner Situation. Diese Wahrnehmung ermöglicht es, einen Kontext des emotionalen und physischen Zustandes zu konstruieren, was als erster Schritt zum Verständnis der Nutzerziele gesehen werden kann. Bei einem Menschen basiert die Wahrnehmung auf Sinnen, so auch bei einem technischen System. Um dies zu erreichen, können eine Vielzahl von Sensoren genutzt werden, die unterschiedlichste physikalische Größen messen. Zum Beispiel ist es möglich die physischen Aktivitäten einer Person, wie bspw. Laufen, mit Hilfe von Beschleunigungssensoren im Mobiltelefon zu messen. Aufbauend auf diesen Daten können dem Nutzer mögliche Reiserouten oder Termine dynamisch angezeigt werden. Das System hat somit einen Teil des Nutzerziels identifiziert und kann dynamisch Unterstützung bieten. Ein anderes Nutzungsszenario, das inzwischen weit verbreitet ist, ist das automatische Entsperren des Smartphones wenn das Gesicht des Nutzers in der Nähe ist. Die Wahl der jeweiligen Sensortechnik, oder auch Modalität, ist sehr schwierig, da unterschiedliche Faktoren miteinander abgewogen werden müssen. Während dieses Prozesses offenbaren alle Technologien eine Reihe von Vor- und Nachteilen, die sehr stark vom endgültigen Nutzungsfall abhängen.

Die Erweiterung der Wahrnehmungsmöglichkeiten von bestehenden Sensoren findet in der Wissenschaft seit jeher große Aufmerksamkeit. Ein weiteres Ziel ist die Reduktion von Hemmnissen während des Designs von Sensorsystemen, wie bspw. die Verringerung des Energieverbrauchs oder des Platzbedarfs. Im ersten Teil dieser Arbeit erörtere ich technologieübergreifend Ansätze, um die Umgebung und einen Nutzer wahrzunehmen. Ich orientiere mich an drei Wahrnehmungszielen: Der Interaktion mit Oberflächen, der Erkennung von Körperbewegungen und der Erkennung von Interaktionen mit Objekten. Mit diesen Zielen betrachte ich unterschiedliche Sensortechnologien, wie Kameras, Ultraschallsensoren, Infrarotsensoren und Ansätze basierend auf elektromagnetischen Wellen. Alle Technologien haben ihre spezifischen Eigenarten, welche unter anderem in der Auflösung und Reichweite, im Energieverbrauch und in der Platzierung liegen. In meiner Arbeit konzentriere ich mich auf die Wahrnehmung eines Nutzers mit Hilfe von quasistatischen elektrischen Feldern oder auch kapazitive Sensorik. Zunächst erläutere ich die Funktionsweise kapazitiver Wahrnehmung im Detail, beginnend bei schwach elektrischen Fischen in der Natur bis hin zu modernen Gestenerkennungssystemen. Auf Basis dieser Technologie ist es möglich, Nutzerschnittstellen zu entwickeln, die in der Lage sind Gesten, Körperbewegungen und Manipulationen der Umgebung in bis zu 50 cm Entfernung zu erkennen. Im Gegensatz zu anderen Methoden, wie bspw. Kameras, ist kapazitive Sensorik nicht abhängig von Beleuchtung und visueller Verdeckung. Aufgrund der geringeren Auflösung haben diese Sensoren einen geringen Einfluss auf die gefühlte Privatsphäre des Nutzers. Kapazitive Sensoren sind sehr energieeffizient und können unauffällig unter allen nichtleitenden Materialien angebracht werden. Basierend auf dieser Modalität wurden Nutzerlokalisierungssysteme [VMV09, SL08], tragbare Aktivitätserkennungssysteme [CAL10, CGL*12] und intelligente Möbel [WKBS07b] realisiert.

Wie bereits zuvor beschrieben, ist eine Wahrnehmung der Umgebung essenziell um die Interaktion mit Technik zu vereinfachen. Auf Basis von Informationen über den Nutzer können sich technische Systeme intelligent anpassen und automatische Entscheidungen treffen. Um dieses Ziel mit Hilfe von kapazitiver Sensorik zu erreichen, orientiere ich mich in den folgenden Kapiteln an drei Forschungsfragen: (1) Der Erweiterung kapazitiver Wahrnehmungsmethoden, (2) dem Erkennen von Körperteilen eines Nutzers auf Basis kapazitiver Sensoren, und (3) dem expliziten und impliziten Interaktionsdesign solcher Systeme.

Die erste Forschungsfrage bearbeite ich mit einer neuen Plattform *OpenCapSense* zur Prototypenentwicklung mit Hilfe von kapazitiver Sensorik [GPBB*13]. OpenCapSense ist in der Lage, Körperteile in 35 cm mit einer Auflösung von ca. 1 cm zu erkennen. Es unterstützt zudem sehr hohe Aktualisierungsraten im Bereich von bis zu einem Kilohertz. Somit lassen sich schnelle Interaktionen, wie bspw. Stürze durch Personen oder durchgeführte Gesten, erfassen. Gegenüber einer bestehenden Plattform bietet OpenCapSense eine signifikant höhere zeitliche und räumliche Auflösung. Ich validiere OpenCapSense mit einigen Anwendungsbeispielen, wie der Sturzerkennung von Personen, einem Gestenerkennungsgerät oder einer interaktiven Kunstinstallation (gezeigt in Abbildung 0.1). Da OpenCapSense nicht in der Lage ist, Menschen und Objekte eindeutig zu identifizieren und sehr feine Manipulationen nicht detektiert werden, stelle ich mich der Forschungsfrage mit einem zweiten Beitrag. Ich untersuche das Konzept der kapazitiven Nahfeldkommunikation (*CapNFC*), um Informationen über Manipulationen und Identität von Objekten mit Hilfe von kapazitiver Kopplung zu übertragen. Die Umgebung eines Objektes kann somit als Informationsraum angesehen werden, in dem Nachrichten veröffentlicht werden können. CapNFC kombiniert Sensorik und Kommunikation in intelligenter Weise. Da der menschliche Körper einen Einfluss auf die kapazitive Übertragung hat, können bspw. Näherungen und Berührungen detektiert werden. Mit Hilfe von CapNFC ist es ebenfalls möglich, Informationen durch den menschlichen Körper zu senden. Insbesondere diese Kommunikationsmethode ist sehr innovativ, um intuitive und höchst interaktive Systeme zu designen. Die beiden wissenschaftlichen Beiträge zu dieser Forschungsfrage werden in den Kapiteln 3 und 4 vorgestellt.

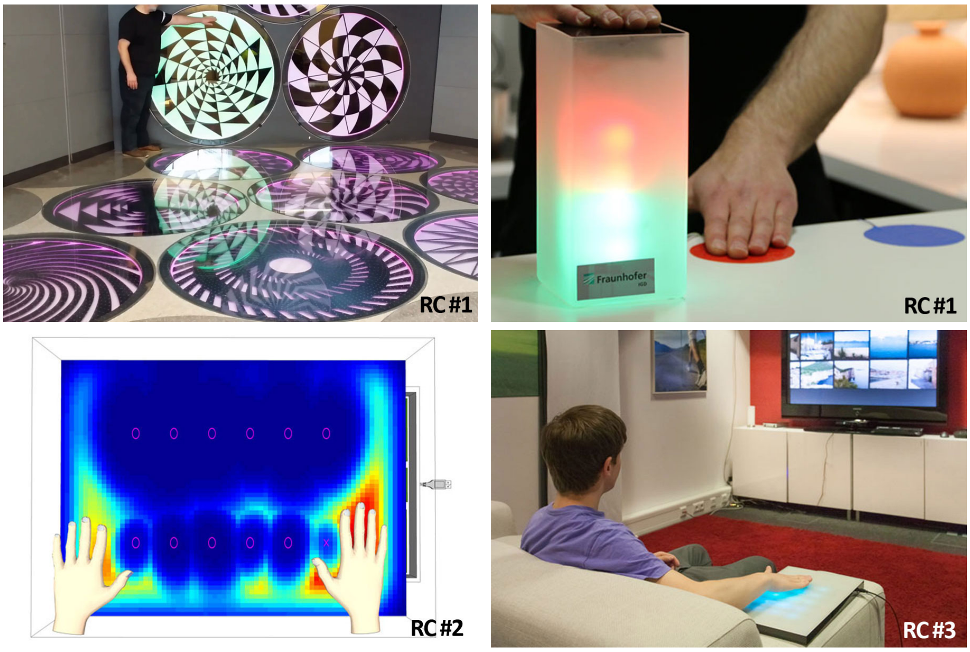

Abbildung 0.1.: Die wissenschaftlichen Beiträge bearbeiten drei Forschungsfragen: (1a) Prototypisierung mit kapazitiver Näherungssensorik, (1b) eine neue Methode zu kapazitiver Nahfeldkommunikation, (2) eine Methode zur Erkennung von Objekten mit kapazitiven Näherungssensoren und (3) der Untersuchung von impliziten und expliziten Interaktionskonzepten

Innerhalb der zweiten Forschungsfrage konzentriere ich mich auf die Erkennung von Körperteilen oder abstrakt Objekten mit Hilfe von kapazitiven Sensoren. Hier existieren viele Ansätze zur diskreten Klassifikation von Sensorwerten, während Verfahren zur Extraktion von kontinuierliche Informationen weniger stark vertreten sind. Um Objekte mit kapazitiven Näherungssensoren zu erkennen, gibt es bereits einige Ansätze, die keine detaillierte Informationen über mehr als drei Freiheitsgerade eines Objektes erkennen können. So gibt es zum Beispiel aktuell kaum Methodik, um neben der 3D-Position das Öffnen und Schließen einer Hand zu erkennen. Mit *Swiss-Cheese Extended* stelle ich einen Verarbeitungsalgorithmus zur Objekterkennung vor [GPBKK13]. Der Algorithmus erkennt mehrere Objekte im Interaktionsbereich über einem Gerät und verfolgt diese während der Interaktion. Swiss-Cheese Extended nutzt einen Partikelfilter, um mehrere Hypothesen effizient und in Echtzeit evaluieren zu können. Die Methode wird mit einem eigens entwickelten kapazitiven Gestenerkennungssystem [GPB12] evaluiert. Der wissenschaftliche Beitrag findet sich in Kapitel 5.

Die dritte Forschungsfrage zum Interaktionsdesign kapazitiver Systeme bearbeite ich zunächst aus impliziter Sicht in Kapitel 6. Hier ist das Ziel die Nutzersituation implizit zu erkennen, ohne dass der Nutzer direkt und willentlich mit einem System interagiert. Die Erkennung dieser Situation trägt somit zum Ausführungskontext des Systems bei und kann zu einer intelligenten Systemreaktion führen. Ein Unterbereich der impliziten Interaktion ist die Erkennung physischer Aktivitäten, welche auch als Aktivitätserkennung bezeichnet wird. Ich untersuche zunächst die Erkennung von Aktivitäten mit körpergetragenen kapazitiven Sensoren, die die Nähe und Natur von Gegenständen messen können [GPBB12]. Im Folgenden wende ich mich stationären Installationen mit Fokus auf intelligente Möbel zu. Ich evaluiere anhand eines augmentierten Sofas und eines Schreibtisches die Erkennung von Nutzerposen und Aktivitäten [GPBK*13,GPMB11]. Im Rahmen der gewollten, also expliziten Interaktion, setze ich mich mit der Nutzerfreundlichkeit kapazitiver Interaktionssystem auseinander. Insbesondere wenn Sensoren unauffällig in der Umgebung platziert werden, ist es für Nutzer häufig nicht ersichtlich, dass eine Interaktion möglich ist. Außerdem ist unklar, wie Interaktion und Systemreaktion zusammenhängen und welche Arten der Interaktion unterstützt werden. Mit Rainbowfish präsentiere ich ein kapazitives Gestenerkennungsgerät, welches mögliche Interaktionen auf der Oberfläche visualisiert [GPBW14a, GPBW*14b]. Somit können Gesten und Systemreaktionen angedeutet werden und dem Nutzer Rückmeldungen zu erfolgten Interaktionen bereitgestellt werden.

Ich schließe meine Dissertation mit einer Zusammenfassung und Diskussion der behandelten Themen ab. Außerdem präsentiere ich einen Ausblick auf Arbeiten, die ich im Rahmen meiner zukünftigen wissenschaftlichen Arbeit weiter untersuchen möchte.

Contents

List of Figures

1. Introduction

The technical functionalities of electronic devices in a typical living environment constantly grew in the last decade. Among the reasons is the fast technological progress in the fields of Electrical Engineering and Computer Science. During this development, existing device types, such as personal computers, were miniaturized and new usage scenarios requiring new devices, like smartwatches, arose. The technological progress over the last couple of decades lead to a vast variety of embedded systems with many different human-machine interfaces. People making use of existing and new technologies experience problems as an ever-growing number of usage patterns has to be internalized. These heterogeneous usage concepts commonly lead to problems in bailing out the capabilities of modern technologies among all generations. Another problem becomes inherent when combining multiple devices to achieve a singular goal. Currently, many devices must be controlled separately while different usage steps have to be executed in the right order.

In this decade, interaction techniques like gesture-recognition systems or speech interaction applications emerged. These systems tackle the common goal of making interaction more easy, natural, and enjoyable. Unfortunately, many of such systems still require training for novice users in advance. However, these more natural forms of human-computer interaction are certainly a step towards the right direction. The way towards more generalizable and comprehensive interaction design still draws great attention in the research community. Besides an increased amount of interoperability between multiple devices, understanding the user and the environment can lead to more adaptive and intelligent interaction approaches. According to a famous article by Mark Weiser [Wei99], the goal is to seamlessly integrate devices in an environment which sense a user's needs in order to give support in achieving her or his goals. This fact also allows the technology to disappear from the user's perception, making the technology a fundamental and indistinguishable part of the environment [Wei99]. However, even 15 years after this vision was formulated, the trend towards simplification and more intelligent user interfaces did not yet pervade the majority of application domains.

One of the key drivers for achieving Weiser's goal is the perception and interpretation of a user's environment. This perception enables to construct a context of a technology user's physical state, which can be regarded as a step towards understanding the user's goals. Perceiving and understanding the user can be realized with a vast variety of sensing approaches. For example, accelerometers are able to sense motions from mobile phones to analyze a person's movements and activities. Based on these activities, information about possible travel routes and appointments can be dynamically adapted. Moreover, modern smart-phones unlock the user interface automatically when the user's face appears in front of the smart-phone's camera. Even though the variety of user interfaces is very high, the choice of the most suitable sensing technology is always crucial. During this decision process, all technologies reveal a number of advantages and disadvantages which depend on the specific use case and its constraints.

This thesis connects two important subjects, *environmental perception* and *ubiquitous interaction*, as depicted in Figure 1.1. Throughout the thesis, the *environment* of a person or device is regarded as a region within its perceivable bounds. These bounds are limited by the person's or device's senses, for example by acoustic or visual perception. Occupancy sensors can act as a simple example: The visual sensor is limited by walls or detection distance, often restricting the sensor's environment to a room. Mobile devices or humans can potentially perceive a subset of the whole world, which leads to a much

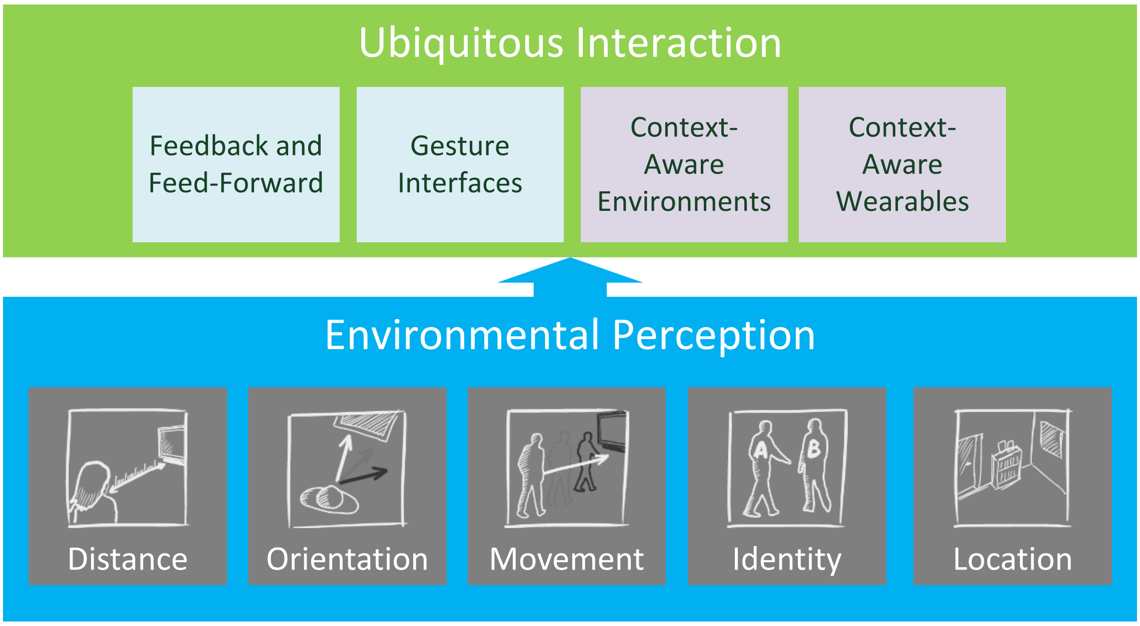

Figure 1.1.: The five proxemic interaction dimensions [GMB*11] are the basis for environmental perception. They enable ubiquitous interaction on an explicit and implicit level.

more dynamic view of the environment. Based on the perception of an environment, devices are able to infer explicit or implicit interactions which enable ubiquitous interaction with humans. For example, they can adjust their user interface to activities being recognized, or react to intentional commands or gestures carried out by a user.

In order to allow for environmental perception, I use the notion of *proxemic interactions* in UbiComp, which was coined by Greenberg et al. [GMB*11]. The authors were inspired by Edward Hall's article on spatial relationships between people [Hal66]. In this work, Hall shaped the term *proxemics* with a division in four human-centric proximity zones: intimate distance, personal distance, social distance, and public distance. Abstracting from the human-centric view enables to project the concept on entities which include people, digital and non-digital devices. Greenberg et al. introduce five proxemic dimensions, that are shown in Figure 1.1. These dimensions provide means of characterizing proxemic interactions in UbiComp. They are not necessarily measurable in continuous scales, they also include discrete observations.

1.1. Motivation

As described in the previous section, a detailed environmental perception is essential for enhancing human-machine interfaces. Data about the user and the surrounding environment can be used to intelligently adapt the interaction possibilities or even make decisions without the necessity of user feedback. Sensing a user's context, which may include the position of body parts, activities carried out, or intentionally given commands, can be achieved by a plethora of different sensor types.

The difficult choice of a suitable sensor technology can be sketched out very easily by the simple example of sensing a user's motions in a room. This goal can be achieved by approaches like cameras, infrared sensors, capacitive sensors, or accelerometers. Each of theses approaches reveals its individual strengths in different usage scenarios. For example, cameras often lack of a high power consumption due to extensive use of processing algorithms. On the other hand, they provide the *perceptional capabilities* to sense very fine-grained body movements. Therefore, it is currently only feasible to use high-resolution cameras in embedded systems which allow for a high *energy consumption*, for example by enabling easy recharging or integrating a stationary power supply. Moreover, in some situations user's may not accept the deployment of cameras even if they are not able to exchange information with other systems, for example in a bath room [Kir14]. Infrared sensors are able to measure distances to objects, with a lack in the perceptional capabilities when direct sunlight is involved. Moreover, they have a limited interaction distance, making it necessary to deploy an array of sensors. The choice of *sensor placements* is also difficult when using energy-saving accelerometers which are usually bound to a moving objects. Capacitive sensors are very energy-saving and can be installed unobtrusively under any kind of non-conductive surface. However, similar to infrared sensors, they also lack of interaction distance. It can be concluded that the choice of the most suitable sensing technology is a significant challenge for developers. Trade-offs including energy consumption, sensor placement, perceived privacy, monetary cost, and the perceptional capabilities have to be accepted.

Extending the perceptional capabilities of existing sensor technologies while lowering design trade-offs has found a lot of interest in ongoing research. Especially in the last couple of years, this approach was very successful in the field of perceiving the environment with quasi-electrostatic fields, or capacitive sensing. By employing this technology, it is possible to implement interfaces that are able to determine gestures, body movements and environmental changes at typical distances up to 50 cm [SGB99]. In contrast to camera-based methods, capacitive sensing has the advantage of being robust against changing lighting conditions and visual occlusion. Additionally, capacitive sensors have a low impact on a user's perceived privacy, while the actual privacy can be very high among all technologies. Energy-efficient capacitive sensors can be deployed unobtrusively underneath furniture, carpets or within walls. The sensed data can be processed with computationally cheap algorithms. The drawbacks of using capacitive sensing are a limited resolution and error-proneness in environments with many conductive objects. Moreover, other electrical devices can affect the sensors' generated electric fields and thus induce high noise. Using capacitive sensors, researchers have realized location tracking systems [VMV09, SL08], wearable activity recognition systems [CAL10, CMPT12], smart furniture [WKBS07b], and gesture recognition systems [WKBS07a, SGB99].

Research is needed to lower the design trade-offs when using capacitive sensing. This modality offers many unique advantages which comprise high speed, unobtrusiveness and interactivity. However, many sensing problems are solved with other technologies that have richer perceptional capabilities. Unfortunately they also consume more power, need more space, induce higher monetary cost, and are less unobtrusive. Therefore, an extension of perceptional capabilities for capacitive sensors is the door-opener for a new generation of both highly-interactive and low-power devices.

1.2. Research Challenges

In the last section I outlined the reasons for research in extending the perceptional capabilities of capacitive sensors. In order to achieve this goal, I identified three research challenges which I target in this dissertation. The challenges start off with the lowest layer, aiming at investigating novel sensing approaches. The second challenge is on the interpretation of data generated by the previously mentioned sensing approaches. Contributions within the previous research challenges enable to investigate the gained interaction possibilities and introduce new interaction paradigms. In the following, I describe the three research challenges in detail.

(1) New capacitive sensing approaches: At first, the choice of a suitable capacitive sensing approach is vital for the performance and perceptional capabilities of the overall system. Current commercial sensing solutions are rather limited in their capabilities as they hide their signal processing from the application developer. Unfortunately, this policy forbids to understand the at firsthand undesired effects in capacitive sensing. For example, sensing noise artifacts may lead to novel ways of distinguishing active electronic devices with switch-mode power supplies. Therefore, the first research challenge is the development of new capacitive sensing approaches in terms of hardware and software. These may pave the way for a better and extended environmental perception as well as the exploitation of new interaction capabilities.

(2) Interpretation and fusion of data from capacitive sensors: Distinguishing touch from non-touch is a rather trivial task whereas sensing a human's body parts in proximity is far more complex. Moreover, in order to extract fine-grained object-related properties, measurements from multiple capacitive sensing sources must be fused in an intelligent way. This task also leads to the problem of fusing data from sensors which are required to function in various geometric constellations. As this approach often depends on specific use-cases, these must be taken into account as well.

(3) Interaction design based on novel perceptional capabilities: The third research challenge focuses on explicit and implicit interaction with the new low-level methods. As motivated previously in this chapter, users often experience problems when using novel human-machine interfaces. This raises the question how an extended perception can contribute towards a better understanding of the user. Furthermore, when interacting explicitly with a system, it is necessary to develop new interaction concepts which are intuitive and natural. Besides sensors, actuators are necessary to companion the pure sensing approach. They can comprise light or sound used to provide natural mappings and introduce usage constraints.

1.3. Contributions

In the following, I describe my scientific contributions to the previously presented research challenges. My contributions target two research fields in computer science: Ubiquitous Computing (UbiComp) and Human-Computer-Interaction (HCI). The contributions shown in Figure 1.2 can be assigned to the three research challenges. The first challenge focuses on new physical sensing opportunities, the second challenge on processing of capacitive sensor data and the third challenge on applying the novel interaction opportunities in terms of new usage concepts and scenarios.

(1) New capacitive sensing approaches: Many current commercial capacitive sensing systems suffer of the capability of fine-grained proximity detection. Moreover, these systems use signal processing approaches which prevent the realization of certain use-cases in which unfiltered data is required. This lead me to the development of a novel open-source toolkit which allows for high-speed acquisition of proximity data [GPBB*13]. It supports a wide range of possible load capacitances, which is an advantage to current commercial systems. The system is evaluated in a number of use cases with a detailed investigation of electrode size, materials and sensing configurations. A second contribution to this research challenge is a generic method for capacitive near-field communication (CapNFC). So far, capacitive communications have been applied for identifying persons at touch screens [VG13], as well as data transfer between touched devices through the human body [YCL*11, Zim96]. However, up to now, no generalized methodology has been proposed. In this contribution I investigate capacitive communications by identifying and evaluating a set of operating modes [GPHW*14]. It paves the way for a low-power communication link between various smart objects and enables new opportunities in interac-

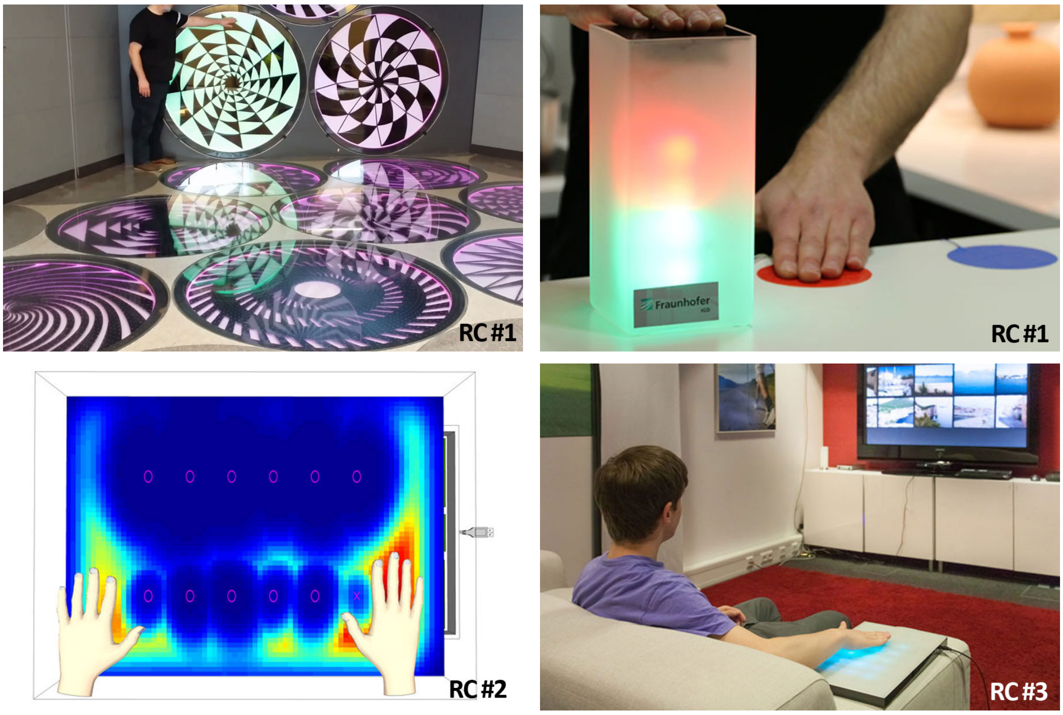

Figure 1.2.: The contributions in this work focus on the three research challenges which comprise (1) new technologies for capacitive sensing, (2) methods and algorithms for capacitive sensing, and (3) new interaction concepts.

tion design. The operating modes are evaluated quantitatively based on a reference implementation of a transceiver and a simple tag.

Interpretation and fusion of data from capacitive sensors: Many current systems focus on recognizing touch with capacitive sensors, while proximity detection is an emerging research topic with first industrial applications. However, when it comes to 3D interaction or sensing the properties of complex objects, there exist only few approaches to determine these object configurations. In this contribution, a generic method is introduced for recognizing continuous object parameters [GPBKK13]. It is applied and evaluated in the field of multi-object body part recognition. Based on such methods, use-cases like gesture recognition in front of displays can be realized [Ber12].

Interaction design based on novel perceptional capabilities: Extending sensing opportunities of capacitive systems induces the need for new interaction concepts. These concepts comprise the fields of explicit and implicit interaction. In implicit interaction, recognizing a person's physical state is vital for understanding the user and support the user's goals. Therefore, I present contributions in the field of activity recognition with wearable sensors [GPBB12] and stationary deployments [GPMB11, GPBK*13]. Considering explicit interaction, current capacitive interaction systems lack of providing natural mappings, usage constraints and ways to provide feedback. By investigating dynamic lighting, I introduce a new feedback and feed-forward modality which enables a more intuitive usage of capacitive interaction systems [GPBW14a, GPBW*14b]. Moreover, by applying capacitive sensing in physical objects, an even stronger mapping to the more obvious physical affordances can be achieved [GPHW*14].

1.4. Structure of this Work

The previous section summarized the scientific contributions made with this thesis. Besides introduction, related work, and conclusion, my dissertation comprises five chapters with scientific contributions. At first, I will give an overview of related work in the field of environmental perception. I introduce how capacitive coupling is used as a sensing organ by animals, namely weakly electric fish. After presenting the history of capacitive sensing, I focus on physical sensing opportunities. They consist of three sensing goals, namely sensing interactions on surfaces, body movements, and object usage. I compare different technologies to capacitive coupling techniques, for example acoustic sensing and cameras.

The first research challenge is tackled by contributions presented in Chapters 3 and 4. First a novel prototyping toolkit for capacitive proximity sensing applications is introduced. *OpenCapSense* allows for prototyping a large variety of capacitive sensing applications, for example fall recognition systems or wearable devices. It is able to provide distance measurements with a resolution of 1 cm at object distances of 35 cm. The sensors can be sampled at rates up to 1 KHz, which enables to design especially high-speed interaction systems. It supports three different measurement modes, making it a versatile and flexible technology for rapidly prototyping ubiquitous interaction systems. However, OpenCapSense is not able to recognize the identity and fine-grained manipulations of objects. In order to complete that picture, I discuss Capacitive Near-Field Communication (CapNFC) as a method for mutual collaboration among objects in the following chapter. Using this method, it is possible to communicate information through air in distances up to 20 cm and through the human body. This allows objects to broadcast information about their acceleration, as well as their ID, to other objects in proximity. The method can be realized with a very low power consumption and represents can be seen as an interactive companion to RFID. Existing objects can be equipped very easily with communication abilities, as CapNFC only requires a single microcontroller pin for unidirectional communications.

A contribution to the next research challenge is presented in Chapter 5. *Swiss-Cheese Extended* is a method for object recognition and tracking based on capacitive sensing. By fusing proximity information from multiple sensors, it is possible to recognize human hands or body parts with multiple degrees of freedom. These degrees can comprise the pitch, yaw, and roll of a human hand. In order to make the method ready for real-time processing, hypothesis about object configurations are approximated over time using particle filtering. The method is evaluated with a custom-built gesture-recognition device.

The next two chapters contain contributions within the third research challenge. Chapter 6 describes the use of capacitive sensors to enhance context-awareness of devices and the environment around a user. The goal is to recognize physical activities which contribute to a common understanding of a user's situation and goals. Here, I am most concerned about implicit interactions that are perceived by a computing system and are not conducted with the intention of interacting with it. I present case studies that exploit sensor placements on the human body and stationary deployments, for example in smart furniture. In Chapter 7, I discuss the use of capacitive sensors in explicit, and thus intentional, interaction scenarios. Recent advancements in technology have led to the disappearance of mechanisms that signal users interaction capabilities. Novice users are often not able to immediately interact with gesture recognition systems, as signaling mechanisms vanish. In order to solve this problem, I introduce the concept of a low-cost gesture-recognizing surface that visually projects information on its surface. This information indicates possible gestural movements, gives feedback and indicates the outcome of actions.

Chapter 8 concludes my thesis with a summary and identifies areas of future research. Last but not least, the appendix comprises my prototype's schematics.

2. Background & Related Work

In this chapter, I will give an overview about state-of-the art sensing approaches to allow for environmental perception. Furthermore, I will provide supporting background knowledge to each technology. The first part of this chapter will be concerned with capacitive coupling approaches, while the second part builds up a context to alternative technologies. Related work which is specific to application domains is presented in each of the following chapters.

Before focusing on technology, I will first describe the goals I would like to achieve with it. I aim at perceiving the environment around a sensing device to sense human interactions. These interactions can happen in very small to larger scales, ranging from sensing interactions on surfaces to sensing whole-body movements. On the one hand, such interactions can be implicit, meaning that the device solely perceives the human body without the user's intention to interact. On the other hand, one would like to capture explicit, and thus intentional, interactions. Perceiving this large bandwidth of possible interactions can be achieved very elegantly by capacitive sensing techniques. Therefore, the first part of this chapter will present background information and related work in this domain. In the latter part, I will present alternative technologies, like cameras, microphones or ultrasonic sensors.

Michahelles and Schiele presented a very nice classification approach to physical sensing opportunities [MS04]. Their work shows that the multi-dimensionality in classifying such physical sensing modalities is highly challenging. In order to solve this problem, Braun developed a benchmarking model for sensing user interactions which includes multiple factors like placement, speed, and quality of sensing [Bra14]. Although there are plenty of sensing technologies and dimensions to consider, the final sensing goals remain the same [MS04]. This led me to the distinction of the following sensing goals: (1) sensing interaction with surfaces, (2) sensing body movements, and (3) sensing object usage.

Aside from the unique technological possibilities, one has to include a number of requirements and constraints for each technology discussed. These include deployment constraints, energy consumption, cost, complexity, and perceptional capabilities. Moreover, soft factors such as privacy awareness and irrational personal fears concerned with electro-magnetic radiation or electric fields must be taken into account.

2.1. Ubiquitous Interaction and Environmental Perception

This thesis addresses two large research areas - human-computer-interaction (HCI) and Ubiquitous Computing (UbiComp). These areas comprise many intersecting concepts and technologies. *Ubiquitous Computing* is the notion of an environment composed of a multitude of computing systems [Wei99]. The technology disappears from the user's perception and becomes an indistinguishable part of the environment. Thinking further towards the vision of *Ambient Intelligence*, the environment supports the user in achieving personal goals [AW09]. As the notion reveals, it requires means of artificial intelligence paired with ambient sensing and actuation.

Explicit interaction is the traditional form of interacting with technology and computers. Here, a user triggers a discrete action and expects a response from the targeted system [WO05]. For example, pressing a button in a graphical user interface should result in some kind of action carried out by the computing system. Triggering such actions is usually achieved by interaction mechanisms which reveal themselves to the user [Nor02]. For example, a button conveys the ability of being pressed, or a door knob transports the ability to be pulled. In HCI, these signaling mechanisms, or *affordances*, are transferred to the virtual world, for example by placing shadows behind areas that can be clicked on. A particular challenge is to convey such affordances with regard to the disappearance of technology envisioned in UbiComp. *Implicit interaction* refers to the concept of perceiving the user, even when direct interaction is not intended [Sch00]. Knowledge about the user's situation enables a system to independently trigger intelligent actions.

In order to sense implicit and explicit interactions, capabilities must be provided to the device or environment that allow for *environmental perception*. Traditional graphical user interfaces (GUIs) have successfully applied mouse and keyboard as input modality. With the ongoing technological progress, new modalities are able to meet the users' needs for easy and natural interaction. This resulted in a trend towards perceptional user interfaces (PUIs) which are based on natural interaction [TR00]. They comprise techniques like gesture recognition or speech interaction. A subset of these natural interaction techniques are *proxemic interactions* [GMB*11]. Proxemic interactions rely on sensing a number of dimensions associated to an object or human. The dimensions are composed of distance, orientation, movement, identity, and location. The obtained information is highly valuable as it can provide means for recognizing both explicit and implict interactions with a computing system.

The previous paragraphs raise the question at which point interactions can be regarded as *ubiquitous*. Clearly, most of the previously introduced aspects of UbiComp should be fulfilled. Explicit interaction should not be targeted towards a specific system's implementation. The user should be able to express the goal by means of natural interaction, let it be speech commands or gestures. Implicit interaction can be regarded as ubiquitous by nature, as the paradigm aims at understanding the user's situation. A further criterion is the involvement of multiple devices in interaction. This requires capabilities in terms of exchanging and understanding information. Furthermore, an important factor to achieve Ubiquitous Interaction is the disappearance of technology in the environment.

2.2. Principles of Capacitive Sensing

2.2.1. Capacitive Coupling in Nature

While human exploitation of capacitive coupling reaches back about one century, nature has been using the effect for millions of years. In regions where the environmental perception is rather limited due to natural constraints, for example due to missing sunlight, some species have developed a very interesting form of exploiting capacitive coupling. For example, weakly electric fish characterize their environment by measuring the distortions of an electric field built up from their head to tail [BLL06]. Tiny electroreceptors on the fish's skin register differences in electric potential, allowing the fish to reconstruct an image of the environment. In Biology this approach is known as active electroperception or active electrolocation.

The electric field is produced by a method called electric organ discharge (EOD), which leads to characteristic waveforms and frequency patterns for different species [Hop81]. Besides using electric fields for locating predators and for orientation, active electroreception can also be used for communication amongst different individuals. Interruptions in the periodically performed EODs are often used in fights for threatening the enemy [Hop81]. Passive electroreception relies on detecting fields that are already present and is common for a greater number of species [ECC13]. The electroreceptive species comprise 41% amphibians and 41% catfishes, such as weakly electric fish [ECC13]. Actively generating varying

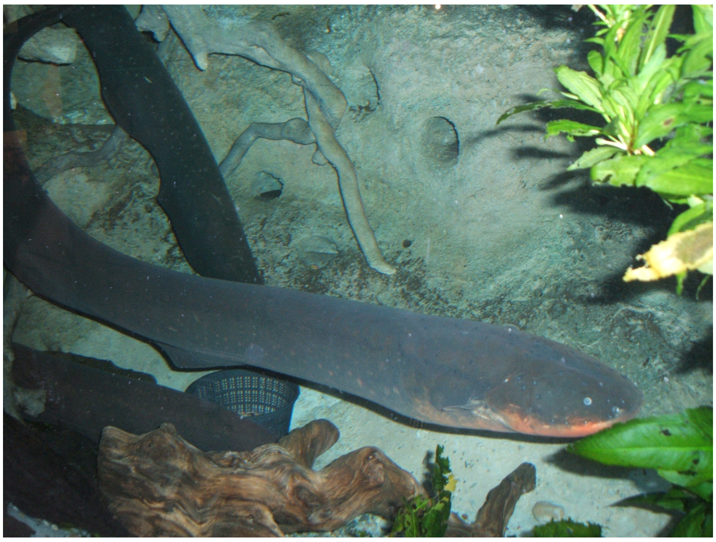

Figure 2.1.: An electric eel which uses active electrolocation. Photo by stevenj (CC BY SA).

electric fields at low frequency and measuring the field's properties is only applied by approximately 60 species [ECC13].

The method of active electrolocation relies on frequency-modulated electric fields with small voltages lower than 1 V with frequencies ranging from 1 Hz to 10 KHz [Nel05]. In some cases, the discharge frequency also maps to the dominance of the individual, where male fish have a lower discharge frequency as female fish [Hop81]. The *Eigenmannia virescens*, a South American fish species, developed a technique to avoid interferences with other fish, known as the jamming avoidance response originally discovered by [WT63]. As soon as one fish interferes with another, both adjust their oscillation frequency depending on the external stimulus. An electric field disturbance leads to a decreasing transdermal voltage on the fish's skin near the object [Nel05, Hei77].

Many principles mentioned above can be mapped to today's technological challenges. For example, the *Eigenmannia virescens* applied a frequency multiplex for interference avoidance long before humans did. Modern touch-screens face the same challenges in detecting objects as weakly electric fish. While electric fish use EODs to achieve both communication and environmental perception [Hop81], the method has not yet been applied in current technology. As the field of bionics shows very prominently, meaningful contributions in science can be made by stepping back and learning from such concepts. One scientific contribution in this thesis is inspired by the method of combining communication and perception using electric fields very much alike electric fish.

2.2.2. History of Capacitive Sensing

Capacitive sensing aims for recognizing objects, such as the proximity to human body parts or conductive materials, using electric fields. Although Leon Theremin [Gli00] is often seen as one of the founding fathers of capacitive sensing, the principle of capacitive sensing was investigated a few years before by the German Cremer [Cre07]. In 1907 Cremer introduced an experiment in which he placed a beating frog heart between two capacitor plates. He measured the change in capacitance with a Saitengalvanometer which allowed him to reconstruct the frog's heart beat without mechanical influence. The Saitengalvanometer was invented some years before by Einthoven, enabling him to measure the first ECG [Pfl95].

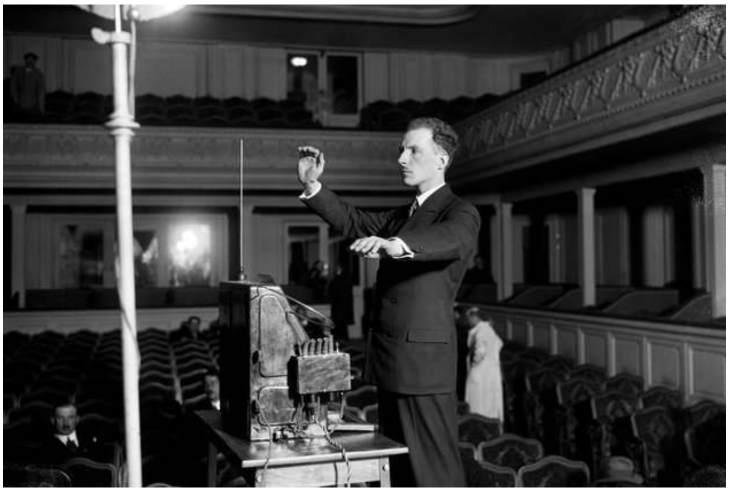

Figure 2.2.: Leon Theremin playing his music instrument based on capacitive sensing. The instrument adjusts its volume and frequency based on the proximity to body parts.

The first practical application of capacitive sensing dates back in 1919 to the Russian physicist Leon Theremin, who invented the first electronic music instrument [Gli00]. Playing his instrument, he is depicted on Figure 2.2. The instrument is composed of two electrodes, both building up an electric field to the surrounding, and thus also human hands in proximity. Even though the hands are mainly used for playing the instruments, the nearby human body has also a significant effect and must be kept calm while playing [PG97]. One electrode is connected to a circuit which influences volume, the other electrode is connected to a circuit for adjusting the instrument's frequency. For each electrode, a resonating inductor-capacitor (LC) circuit was driven at a frequency of around 500 KHz [PG97].

2.2.3. Physical Background

In order to apply capacitive coupling for perceiving the environment it is necessary to understand the fundamentals of capacitive coupling. Throughout the whole thesis, I consider quasi-electrostatic fields. This means that the wavelength is much larger than the largest electrode length applied in our physical setups. Regarding an electric field oscillating at a frequency of up to 1 MHz, the wavelength of 300 m is considerably large for a quasi-electrostatic viewpoint. That also implies that any effects like radiating electro-magnetic fields do not apply in this regime. As the thesis focuses on exploiting capacitive coupling methods in UbiComp and HCI, I will not cover the field of electro-statistics and finding exact solutions to the given problems. Instead, I quickly move to approximating models and their suitability in modeling interactions with capacitive coupling.

A capacitor can be defined as two electrodes which are separated by a dielectric material [CW05]. An electric charge Q can be stored on the capacitor's plates, whereas the amount of charge depends on the capacitance C and voltage V applied to the capacitor. Therefore, each plate of the capacitor either holds a charge of $+Q$ or $-Q$. In order to move the electric charge to the capacitor, a certain amount of work W is required.

$$C = \frac{Q}{V}; W = \frac{1}{2}CV^2 = \frac{1}{2}QV \tag{2.1}$$

2.2.3.1. Displacement Current

When a voltage is applied to the capacitor's terminals, a current is able to flow [Bax96]. The capacitor acts like an open circuit when direct voltages are used, while currents are able to flow using alternating voltages. The total current through the capacitor comprises two components - an induction current I_i and a displacement current I_d. Conduction current corresponds to actually moving charges, a component which is usually very small in capacitors. The much larger displacement current I_d is induced by a changing electric field E through an area A, which corresponds to a charge Q. By differentiating one can obtain the rate of change in electric flux, which is called displacement current I_d.

$$Q = \varepsilon_0 \oiint_A E dA \tag{2.2}$$

$$\frac{dQ}{dt} = I_d = \varepsilon_0 \frac{d}{dt} \oiint_A E dA \tag{2.3}$$

Measuring electric displacement currents is an elegant way to determine the capacitance between two capacitor plates. The method can be applied at both capacitor's terminals, while an excitation voltage is usually applied at one terminal and the displacement current is measured at the opposite one [Bax96]. Changes in the displacement current's amplitude and phase can be applied to reconstruct the capacitance between both electrodes. Even if the phase is unknown, the amplitude still represents an important measure.

2.2.3.2. Dipole Model

In order to obtain an impression about the electric field between conductive plates, the Laplace equation with an inhomogeneous permittivity has to be solved. As this model is very complex and not suitable for real-time calculations, a dipole model with point charges can be applied to model the electric field strength [Smi96].

Although one is usually not concerned with point charges in capacitive sensing, the method of image charges can be applied as means to investigate the electric field with planar or spherical surfaces. However, this method is only valid for homogeneous surface charge densities, which is often not the case in real-world applications [Smi96]. In the dipole model, two point charges are assumed which build up an electric field E [CW05]. The point charge $-q$ is displaced from an opposite point charge $+q$ by distance d. Assuming both charges to be located on the dipole axis $(d/2,0,0)$ and $(-d/2,0,0)$ leads to a dipole moment $\vec{p}$ between both charges.

$$\vec{p} = q \cdot \vec{d} = q \cdot \begin{pmatrix} d & 0 & 0 \end{pmatrix}^T \tag{2.4}$$

For each point given as $\vec{r}$ it is possible to determine the electric field E [CW05].

$$E(\vec{r}) = \frac{1}{4\pi\varepsilon_0\varepsilon_r r^3} \cdot \left(\frac{3(\vec{p}\cdot\vec{r})\vec{r}}{r^2} - \vec{p} \right) \tag{2.5}$$

Based on the equation, one can make two important inferences. First, the greatest electric field is in the center between the both point charges. Second, the field is symmetrical, meaning that there is no difference in absolute electric field strength at a dedicated point when both charges would be turned around. The dipole model has been used to model the relationship between displacement current and the electric field strength [Smi96]. Therefore a *unit absorber* is assumed which can be compared to a very small grounded object. The approach by [Smi96] assumes that the displacement current stands in relation with the electric field strength with the location $\vec{x}$ of the unit absorber. The displacement current is considered proportional to the electric field strength in the given point with an offset K_0: $I_d \propto K_0 - |E(\vec{x})|$ [Smi96].

2.2.3.3. Plate Capacitor Model

A very simple model explaining the relation between two sensing plates, isolator, and the resulting capacitance is the plate capacitor model shown in 2.3 [Bax96].

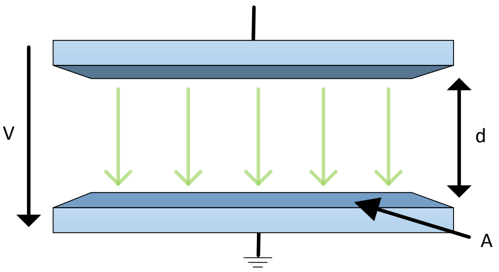

Figure 2.3.: The plate capacitor model is a convenient simplification to approximate the properties of capacitive sensing applications.

In this model, it is assumed that two plates, both having an area of A, are placed parallel to each other with a distance d. The two conductive plates are isolated by a dielectric material, having a permittivity of ε_r with a dielectric constant of ε_0. The resulting capacitance between the two conductive plates depends on A and d.

$$C = \varepsilon_0 \varepsilon_r \frac{A}{d}; \text{and thus } C \propto \frac{1}{d} \tag{2.6}$$

2.2.3.4. Capacitor Charging Time

Inferring a capacitor's size can be achieved by analyzing the charging time of a capacitor. The method of determining a capacitance using an additional resistor is a very common technique in capacitive sensing [Bax96]. In order to determine the time required for charging a capacitor, it is necessary to limit the available charging current. This can be achieved by setting a resistor in series with the variable capacitor, as depicted in Figure 2.4. Due to the current limitation, the voltage slowly increases at the capacitor. When charging the capacitor with capacitance C through a resistor with resistance R, the capacitor's voltage V_c increases corresponding to time constant τ [CW05].

$$V_c(t) = V_0(1 - e^{\frac{-t}{\tau}}) \text{ with } \tau = RC \tag{2.7}$$

Here, τ represents the time which is needed to charge an empty capacitor to 63.2% of full charge and - vice-versa - uncharge a fully charged capacitor to 36.8% of full charge. Measuring charging times is a very convenient way to determine a capacitor's value because it allows for applying easy and cost-effective measurement frontends, such as timers or simple comparators. Usually, this technique is only applied at a single capacitor terminal, without requiring access to the opposite one. Figure 2.4 depicts an example of charging three different capacitors through a resistor R. Measuring capacitances in the region around 15 pF is typical in capacitive sensing scenarios.

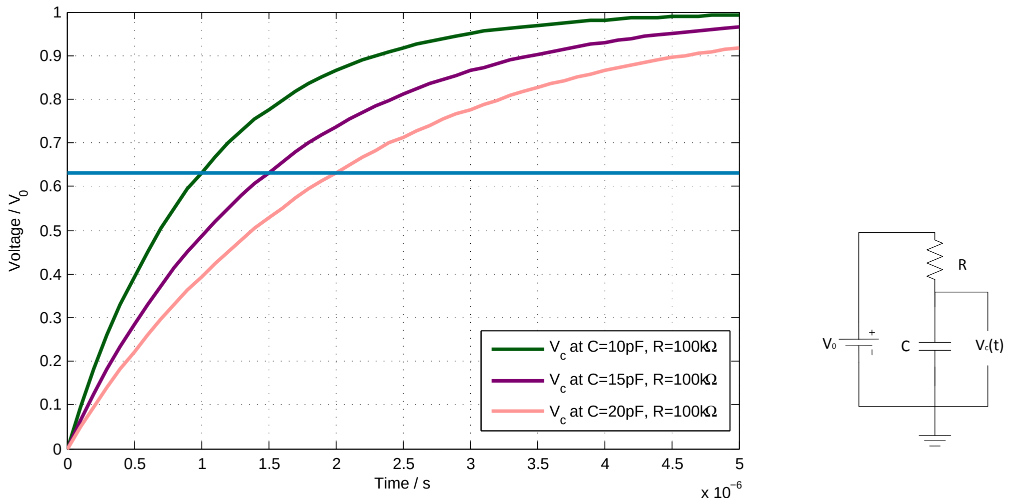

Figure 2.4.: In order to determine the capacitance of C, the time constant can be measured using a limited charging current. V_c depends on the resistance R and the size of the capacitor C. Supposing a fixed value of R and V_0, it is possible to determine the capacitance by measuring the voltage V_c.

2.2.4. Capacitive Sensing Model

2.2.4.1. Generalized Lumped Circuit Model

When applying capacitive sensing in a user's environment, one has to abstract from the view of having singular electrodes. Instead, the whole environment, including human body parts, furniture or a floor can be regarded as parts of the capacitor's electrodes. Such potential environmental electrodes are made of different materials with varying levels of conductivity and coupling to a certain potential. Especially a human body has a very good conductivity, as it is mostly composed of water. Moreover, it has a good coupling to the environment's ground. AC currents can easily flow within the human body, making it an ideal environmental electrode. Other materials such as carpets or wood have a lower conductivity, and thus making them less easier to detect with capacitance measurements. The lumped circuit model by Smith [SGB99] (depicted in Figure 2.5) explains the relation between the different environmental capacitors in a standard capacitive sensing configuration very nicely. In this scenario, the goal is to measure the distance to a human body part H. It shows two electrodes T (transmit) and an optional electrode R (receive). The two electrodes build up an electric field to the surrounding, as well as to the human hand, resulting in capacitances C_{0-4}. The capacitance C_5 is introduced by the coupling of a person to the environment's common ground. For example C_5 is influenced by the person's shoe soles. R_i and C_i represent the corresponding human body's internal capacitance and resistance. As the body is a very good conductor, R_i is in the region below $1\text{k}\Omega$. A detailed description of the human's electric properties will be given in the following section. Both capacitors C_3 and C_4 represent so-called *parasitic capacitances*. In contrast to C_1 and C_2, they are of minor interest for an application developer. As such parasitic capacitances are static, they also affect the sensitivity of a capacitive sensing system. High parasitic capacitance lead to a smaller relative change in measurable capacitance and thus a decreased overall sensitivity. It can be introduced by any conductive object in the region of the sensing electrodes (e.g. metal planes) or caused by hardware components (e.g. input capacitance). Reducing these unwanted capacitances in hardware designs is therefore very desirable and of interest for any application developer.

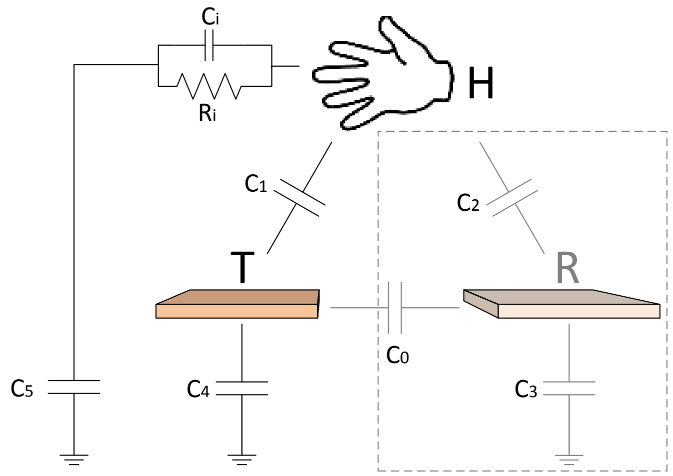

Figure 2.5.: The lumped circuit model by [Smi96] can be applied for both self-capacitance (without electrode R) and mutual-capacitance measurements.

2.2.4.2. Lumped Circuit Model for Self-Capacitance Measurements

First, a sensing technique without a distinct receive electrode R, also known as a *self-capacitance measurement*, is described. Therefore, the electrode T builds up an electric fields to parts in the environment. Both the environment and the human hand act as the opposite electrode to electrode T. The resulting capacitance comprises the two capacitances C_1 and C_4 in parallel. When C_4 can be considered as static, meaning that the environment does not change over time, it is possible to extract the capacitance to the human hand C_1. Moving the hand closer to T increases the capacitance C_1 according to the plate capacitor model. However, determining the distance to the human hand requires strong presumptions about its geometry. It is possible that larger hands or multiple objects lead to the same capacitance as a small hand located at a larger distance. The self-capacitance measurement is very easy to conduct as only a single electrode is required. As described in the previous section 2.2.3.4, measurements based on RC time constants are very applicable for determining the value of C_4. This sensing technique is often applied in comparatively simple capacitive touch buttons.

A sensor for measuring self-capacitance typically uses only a single electrode. Therefore, they are very easy to shield against external influences [VM92]. The principle of applying a *driven shield* is commonly used to reduce parasitic capacitances from electrodes as well as shield the sensor against influences from certain directions. This is depicted in Figure 2.6, where an optional shield is applied with capacitors C_{TS} and C_{4S}. The shield basically represents a second electrode which is driven at the same potential as the sensing electrode above. The driven shield also reduces the capacitance C_4, increasing the potential relative change in capacitance when a hand is brought into proximity of T. Bringing an object close to the shield does not affect the measurement strongly, but leads to an increase of C_{4S}. Applying a driven shield results in very small voltage difference between sensing electrode and shield, which reduces the capacitance C_{TS}. Although the theoretical value of this capacitance C_{TS} should be zero, electronic components can induce small phase shifts or offset voltages.

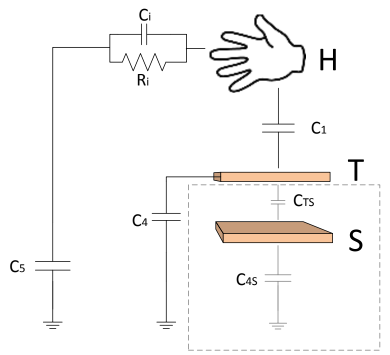

Figure 2.6.: A driven shield can be used to improve the quality of self-capacitance measurements by reducing parasitic capacitances.

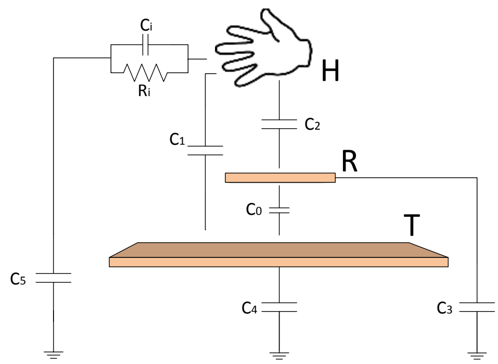

Figure 2.7.: In mutual-capacitance sensing, the transmit electrode can be used for shielding the measurement against undesired influences.

2.2.4.3. Lumped Circuit Model for Mutual-Capacitance Measurements

In order to perform more sophisticated measurements, the model can be extended by a receive electrode R, which is called *mutual-capacitance measurement* (depicted in Figure 2.7). An alternating excitation voltage is applied to T while the incoming displacement current is measured at R. The resulting capacitance between transmit and receive electrode is represented by C_0. A number of additional capacitors must be taken into account: A displacement current also flows from transmit electrode to the human hand and other parts in the environment. When the hand comes close to both electrodes, the displacement current at R decreases. Therefore, the capacitance of C_0 represents a suitable measure of proximity to the human hand. Again, ambiguity considering the size of the target object comes into place, making it difficult to determine an exact distance. Using the sensing technique, one can exploit the combination of multiple transmit and receive electrodes. This qualifies mutual-capacitance measurements to one of the most prominent methods in commercial touchscreens [BO10].

Shielding mutual capacitance measurements is not always as easy as shielding approaches based on measuring self-capacitance. Considering a scenario in which influences from underneath the circuit board should be avoided, it is common practice to route the transmit electrodes on the bottom of the board [Mic14a]. Figure 2.7 shows an exemplary setup, in which the transmit electrode T is used to shield the receive electrode R from the bottom. Unfortunately, this approach unavoidably increases the static capacitance C_0 from transmit to receive electrode and reduces the apparatus' sensitivity.

2.2.5. Capacitive Proximity Sensing

One of the first inventions that made use of capacitive proximity sensing was the Theremin [Gli00]. It was invented in 1919 by the same-named Russian physicist Léon Theremin, who lay the foundations for electronic music and synthesizer techniques [Gli00]. With the rising popularity of information technology and the demand for better and easier human-machine interaction, capacitive proximity sensing became a popular research subject at MIT [ZSP*95, Smi96]. The scientific contributions include

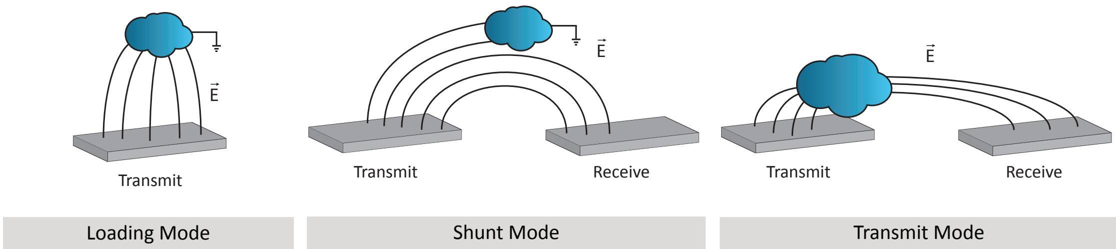

Figure 2.8.: Measurement modes in capacitive proximity sensing identified by Smith [SGB99].

devices that are able to detect hand motion [ZSP*95, Smi96], or more general research about the recognition of human body parts in proximity [SGB99]. Here, Zimmerman et al. and Smith presented the first generalized approaches to proximity sensing. In one of the first works on this topic, Zimmerman et al. presented the *shunt mode* and *transmit mode*. Smith later added the *loading mode*, which is very similar to a self-capacitance measurement. These measurement modes again can be distinguished into either self-capacitance and mutual-capacitance measurements. Mutual-capacitance is often concerned with measuring capacitive properties between two electrodes, while self-capacitance measurements only employs a single transmit electrode. However, it is still possible to mix all modes to achieve better resolutions, decrease detection errors or acquire additional environmental data. In the following I will describe the corresponding modes which are depicted in Figure 2.8.

Loading mode only employs a transmitter electrode. When a human body part approaches this electrode, the capacitance between the transmit electrode and the body part increases. This mode can be regarded as a self-capacitance measurement, in which the same principles apply as described in the previous section. Since only a single electrode is employed in this setup, it is very easy to shield the electrode from external influences. These influences can be body parts approaching from undesired directions, for example from the bottom of the electrode. In order to achieve a good shielding on can either employ a conductive surface with a static potential or use a driven shield [VM92]. In the latter approach, the shield is driven at the same potential as the sensing electrode which leads to a voltage difference which is nearly zero between both plates. This technique reduces the undesired parasitic capacitance, for example induced by an underlying surface. Due to its simplicity, many applications in Ubiquitous Computing use loading-mode sensing. Examples range from capacitive tables and shelves [WKBS07b], to chairs [SGB99] and person-detecting floor systems [Bra09].

In *shunt mode*, a distinct transmitter and receiver electrode is used. The transmitter applies an alternating voltage to the electrode, typically in regions up to 1 MHz. When a human body part approaches the electrodes, the human body shunts the displacement current flowing from transmitter to receiver to ground. Therefore, the measurable displacement current at the receiver electrode is reduced. In most cases the current is measured at the receiver [Smi96], while it can also be measured at the transmitter electrode or at both electrodes. Shunt mode is very applicable for applications in which a high resolution is required. For example, gesture recognition applications [BH09, GPB12, SGB99] benefit a lot from shunt mode sensing, since $\frac{n\cdot(n-1)}{2}$ measurements can be conducted for only n electrodes. When a human body part approaches the transmit electrode in far less than a dipole spacing, the body part itself turns to a transmitter [Smi96]. This results in an increased displacement current for very small body part distances, turning the human body itself to the transmit electrode. The effect refers strongly to the conductive properties of a human body, as described in Section 2.2.7.1. It can be exploited in intrabody communications [Zim96].

Although the effect is undesirable in proximity sensing, using the human body as a transmitter opens new perspectives in interaction design. This method is called *transmit mode* and can be achieved by placing a transmit electrode somewhere near the human body. Possible placements are conductive plates below feet [Zim96] or the application below chairs [DL01]. Very similar to shunt mode, an alternating voltage is applied to the transmit electrode and thus also to the human body. As soon as a human body part comes close to a receiver electrode, the displacement current increases. When a receiver electrode is involved, the approach can be compared to a mutual-capacitance measurement. It is also possible to use the human body as a singular transmit electrode, which can be compared to a loading-mode, or self-capacitance, measurement.

Even though it is useful to classify different operating modes in the domain capacitive proximity sensing, all modes are still tightly connected to each other. They can be combined and used in various cross-over combinations.

2.2.6. Capacitive Intrabody Communication

Depending on the permittivity, a capacitive coupling can be established through different materials. This enables using different materials for sensing or even for information exchange. Most obviously, capacitive coupling can be established through the air. However, coupling is also possible in materials with higher permittivity, such as the human body. Besides transmit-mode proximity sensing, discussed in the previous subsection, it is possible to use intrabody coupling for communications. Here, the human body is used as a medium for communication among objects, as first described in [Zim96]. Unfortunately, problems occur when no common ground between the objects can be provided. This is often the case when battery-powered devices are employed which become less sensitive. In contrast, mutual-capacitance measurements can be conducted easily without having a common ground. In this case, the permittivity of the approaching object plays an greater role than its ability to shunt current to ground. The permittivity depends on the various types of tissue within the human body.

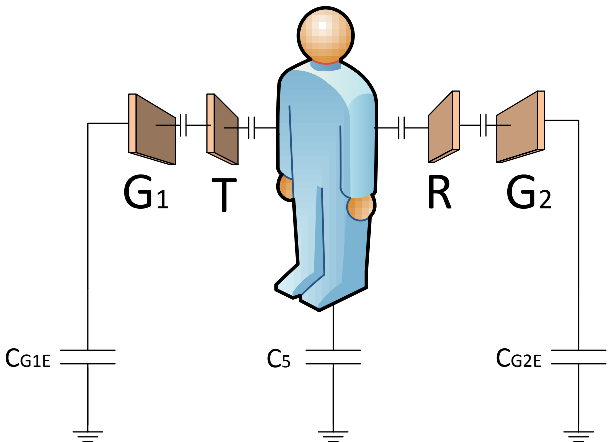

Figure 2.9.: Intrabody communication as introduced by Zimmerman [Zim96]. The capacitors C_{G1E} and C_{G2E} provide the current's return path.

In capacitive-coupled communications, a transmit electrode T sends a message to a receiver electrode R. Zimmerman et al. [Zim96] used the term *intrabody communications* to characterize such methods. As the human body is just a singular medium, the choice of the displacement current's return path is crucial and leads to some of the many limitations of intrabody communications when using battery-powered devices. In order to enable a displacement current to flow, a common ground has to be provided between two or more communicating nodes which are conceptually depicted in Figure 2.17. However, this requires an additional medium besides the human body. Zimmerman solves the problem by attaching additional ground electrodes $G1$ and $G2$ to the nodes, building up an electric field to the environment's ground through air [Zim96]. The resulting capacitance of C_{G1E} and C_{G2E} is only 10 fF small and thus only enables a tiny displacement current to flow between T and R. When using devices connected to a stationary power supply, a common ground is usually available and enables significantly easier intrabody communications.

Based on the human body's conductive properties, it also possible to distinguish between different users on a touch screen [HSP12]. Unfortunately, this interesting property can not act a biometric fingerprint as changing shoes or sweating change the measurable body impedance. This is also the case when the person touches an object, for example by leaning on a table, and therefore changes its connection to the environment's ground.

2.2.7. Capacitive Coupling and the Human Body

2.2.7.1. Conductive Properties of the Human Body

Generally spoken, the human body can be regarded as a very good conductor, as it is mostly composed of water. Various works determined the human body resistance from a hand to the feet to be approximately 100 - 251 Ω [Zim96, Web10]. Considering applications based on capacitive coupling, the low resistivity does not only affect the permittivity ε_r but also enables to shunt a significant amount of displacement current to ground. However, this fact only applies when the capacitive sensing device and the human body share a common ground. This is not the case when capacitive sensing devices run on a battery supply. Using the human body as a conductor also enables new ways of applying capacitive methods. For example, the body itself can be used as an electrode for self-capacitance or even mutual-capacitance measurements.

Averaging the human body's conductivity does not completely make its point, because skin, tissue and bones lead to very different conduction properties. Dry skin for example has a very bad conductivity of approximately 1 $M\Omega$, which suppresses large DC currents to pass through the human body [SPH12]. Nevertheless, AC currents can flow very easily as the human skin and the underlying highly conductive tissue build the opposite plate of a capacitor with the skin as separating dielectric material. Due to the different levels of conductivity, displacement currents take different paths through the human body, which depend on the frequency applied [SPH12]. Frequency sweeps can be exploited to generate more discriminative information about the type of human interaction. Such methods enable to recognize how doorknobs are touched or how arms are placed on a table, using a single capacitive sensor [SPH12]. Figure 2.10 [FL96] shows various types of materials within the human body. While the permittivity of fat and resistance is not dependent on the applied AC current's frequency, the permittivity of muscles and bones differs.

	10 kHz		**1 MHz**	
	Permittivity ε_r	**Resistivity ρ (Ωm)**	**Permittivity ε_r**	**Resistivity ρ (Ωm)**
Bone	640	100	87	50
Fat	30000	15 - 50	NA	15 - 50
Blood	2800	1.5	2000	1.5
Muscle (perpendicular to fibers)	70000	10	1900 - 2500 [1]	1.3 - 1.7
Muscle (parallel to fibers)	80000	2	1900 - 2500 [2]	0.6 - 0.8

Figure 2.10.: Human body impedance of different tissue types and frequencies [FL96].

2.2.7.2. Validity of the Plate Capacitor Model

The simple plate capacitor model can be used very easily to obtain a first impression about the resulting capacitance between a conductive plate and a human body part. The ideal model primarily applies for *self-capacitance measurements*, as only two electrodes are involved in this setup. Although the plate capacitor model is very convenient to apply, one has to ask the question about its validity. This question will be investigated in the following.

The first simplification is that an equally sized human hand and copper plate are considered. In an exemplary setup, this results in a capacitor plate size of $A = 0.1m \cdot 0.1m$. The two electrodes, hand and copper plate, are isolated with air having a permittivity of $\varepsilon_r = 1$. Assuming a hand distance of $d = 0.1m$ results in a capacitance of $C = 0.885pF$, according to Equation 2.6. The results from the ideal plate capacitor model are compared to measurements from a similar real-world experiment in Figure 2.11. As stated above we assume two electrodes, a human hand and a copper plate, having a size $A = 0.01m^2$ and variable hand distances d. The experiment was conducted with a shielded copper electrode placed on a wooden desk. The parasitic capacitance was measured without a body part in proximity and subtracted from the measured capacitance.

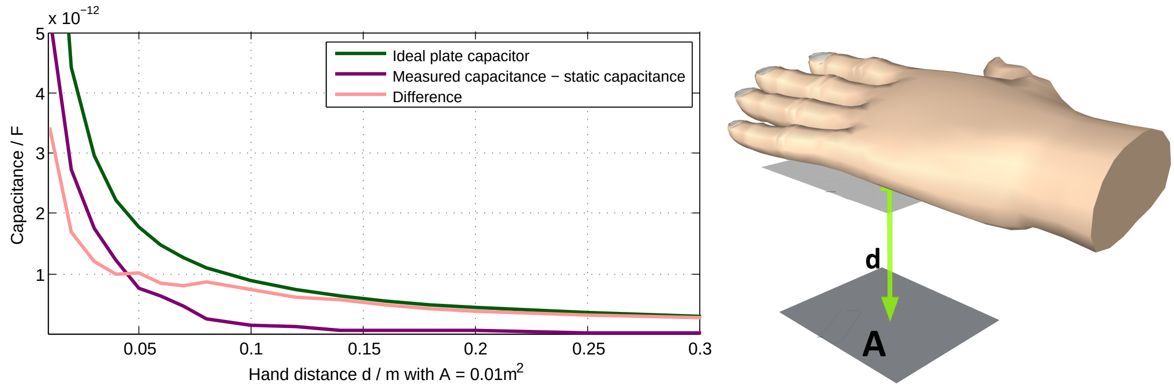

Figure 2.11.: The plate capacitor model compared to a real-life measurement setup (left). The human hand is modeled as the opposite side of the plate capacitor (right) [Ber12].

Especially at hand distances below 0.04 m, the ideal model and the experimental results differ significantly. The reasons for the delta between the ideal model and the experiment are multi-fold. Most importantly, I did not consider parasitic capacitances in the ideal plate capacitor model. These are present in the experimental setup as the electrode was placed on a wooden table and the electronic components themselves introduce parasitic capacitances. Furthermore, the conductivity of a human hand is good but its ability to shunt displacement current to ground is not comparable with an electrode made of copper directly to the environment's ground. In conclusion, the ideal model is a valid way to obtain a first impression about the properties of capacitive coupling applications and their interaction with the human body. However, many properties such as limited electrode conductivity and parasitic capacitances lead to offsets and different behaviour.

2.2.7.3. Validity of the Dipole Approximation

The dipole approximation proposed by Smith [Smi96] represents a very basic way to explain the effects of a human body part approaching to two electrodes. Therefore, the model is mostly suited for *mutual-capacitance measurements*. As depicted in Figure 2.12, the dipole model assumes that objects, such as a unit absorber, simply attenuate the electric field lines in a specific location. This assumption is based on many simplifications. On the one hand, the approximation does not consider that the third object holds a certain charge and transforms the electric field lines between both point charges. On the other hand, objects can usually not be represented by a unit absorber, they are rather volumetric objects.

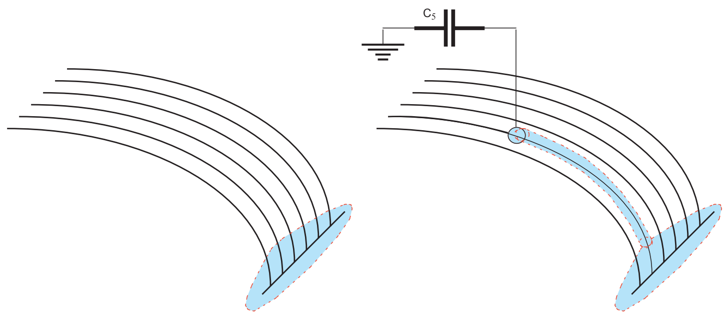

Figure 2.12.: In the dipole approximation presented in [Smi96], the displacement current is proportional to the electric field strength at the location of a unit absorber.

Based on these considerations, one has to ask the questions how these simplifications affect the validity of the approximation and which constraints the it brings along. A very strong constraint is an axis-dependent directivity introduced by the use of planar plate electrodes [GP12]. This means that the measured displacement current depends strongly on the axis on which an object approaches. Spherical electrodes would generate homogeneous field lines, however they are not common in capacitive sensing applications. Although the sensor values should be proportional to $\frac{1}{d^3}$, with hand distance d, measurements have shown that the factor is rather less than 3 [GP12]. This leads to the conclusion that the dipole model provides a basis for modelling sensor responses but needs further adaptations which depend on the actual sensing setup.

2.3. Capacitive Sensing for Environmental Perception

In this section, I present three approaches to allow for perceiving the environment with capacitive coupling. Applying capacitive coupling in the domain of environmental perception brings along advantages and disadvantages. An important aspect is the deployment of capacitive coupling systems: Electrodes and sensors can be placed unobtrusively underneath any kind of non-conductive surface. In terms of detection distance though, they have a rather limited range up to 50 cm. This also represents a drawback for technology developers, which are bound to the physical restrictions of the sensing method. The use of capacitive coupling often results in ambiguous sensor readings which depend on the target object's size and conductivity. This forces developers to cope with false interpretations and induces them to pose strong initial assumptions. In the following, I present approaches for environmental perception in the domains of sensing interaction with surfaces, body movements, and object usage. After all it is not possible to make clear distinctions between the three fields as they overlap in various aspects.

2.3.1. Interactions on Surfaces

Though capacitive sensing is applied in various application domains, touch and grasp sensing is probably one of the oldest applications for capacitive sensing. Originally developed to sense object proximity, recognizing touches and grasps is one of the easiest and most convenient use of capacitive sensing. A very prominent technology for touch sensing are capacitive touch buttons, operating in either self-capacitance or mutual-capacitance mode [Bax96]. They have been applied since decades to sense touches and grasps in application domains like modern touch-screens [BO10], or more trivial use-cases like traffic light switches for pedestrians.

Figure 2.13.: Different electrode layouts can be used to enable a variety of interaction techniques in touch sensing [Bax96, Pra14].

During the last couple of years, some standard electrode layouts evolved to a common concept in capacitive touch sensing. One can divide these electrode layouts into the groups of *buttons*, *sliders*, *wheels*, and *grids* - depicted in Figure 2.13 [Bax96, Pra14]. These different electrode layouts allow for detecting touch-based interaction, such as swipe or rotation gestures. All electrode layouts can be used with either self-capacitance or mutual capacitance sensing. As described in Subsections 2.2.4.3 and 2.2.4.2, the electrodes can also be shielded to the bottom which makes it necessary to include a second electrode layer.

Even though these very common sensing approaches have proven to be stable and support many application domains, research has moved away from standard electrode layouts to application-specific electrode structures. In this domain, rapid prototyping techniques enable users to directly design their favorite user interface on a computer and use devices like vinyl cutters for production [SZH12]. Such examples show that flexibility is a major issue, leading to approaches that investigate the deployment of touch sensors with conductive tape on shaped surfaces [HV11]. Conductive tape can also be combined with other electronic components like LEDs, which can be applied as a sticker [HVC*14]. Inkjet printed circuits and electrodes [KHC*13] represent further important steps towards more flexible layouts of capacitive touch sensors. Olberding et al. introduce cuttable multi-touch sensors that are fault tolerant even when certain conductive lines or areas are being cut [OGT*13].

Besides flexibility in application design, the next major issue in touch-sensing is the deformability of touch surfaces. Current capacitive touch sensors usually act in a very static way. However, flexible and bendable surfaces which allow for touch interaction are very promising technologies in terms of intuitive and ubiquitous usage. Gong et al. [GSO*14] use printed electrode layouts on flexible substrates. Using this approach, it is possible to detect bending and folding of the substrate's edges as well as simple touches on its surface. It can be integrated ubiquitously to make everyday objects like water bottles interactive. Smart garments and textiles can also act as a very flexible substrate for capacitive touch sensors. For example, Cheng et al. apply self-capacitance measurements in textiles to build seamless touch interfaces [CBL08]. Applying these approaches on textiles poses various challenges to the developer. In particular, the sensors must be shielded and undesired interactions due to body movements must be reduced as far as possible [CBL08].

Such works also show that touch sensing is not only applicable in the domain of explicit interaction, but also supports a variety of implicit interaction techniques. For example, Wimmer et al. use touch sensing to discriminate different ways of holding and grasping a tangible device [Wim11]. The device is not only capable of measuring grasps but can also detect desired gestures for a richer interaction with mobile phones.

Touché [SPH12] takes an important step towards more ubiquitous usages and enhanced expressiveness of self-capacitance sensing. By applying a resonant tuning technique, [PYGH10] the conductive properties of the human body are exploited to allow for a richer interaction experience. The presented technique applies a frequency sweep from 1 KHz to 3.5 Mhz using an LC oscillator. The variable capacitance is built up between an ordinary electrode and the touching body part. The system benefits of the changing human body impedance for different frequencies. It enables to detect different types grasps [PYGH10, SPH12], interactions with liquids [SPH12], or the way we place hands on a table [SPH12]. A more visionary instance of resonant tuning can also be found in the Botanicus Interacticus [PSLS12], a sensor-augmented plant. Besides using dedicated sensing electrodes, recent work also integrates existing conductive objects like door knobs [SPH12].

Referring back to a wider view on touch and grasp sensing, not only substrates can act as a deformable structure on which sensing electrodes can be placed. Electrodes like conductive threads can be integrated into smart garments to recognize human activities. For example, Cheng et al. [CAL10] use mutual-capacitance sensors to measure muscle contractions or tissue changes, as depicted in Figure 2.14. As soon as a muscle contraction takes place, the tissue's permittivity changes and leads to a change in electric field between a transmitter and receiver electrode. By placing the smart garment around a person's neck, the authors are able to detect different types of swalling activities like eating and drinking.

Figure 2.14.: Cheng et al. [CABL13] use a capacitive neckband to measure changes in tissue, for example while eating (left). The authors of [CBL08] use capacitive sensing to create touch interfaces on garments (right).
[CBL08] Reprinted by permission of the authors.
[CABL13] © 2013 IEEE.

2.3.2. Body Movements

The use of capacitive sensing to detect interactions on surfaces is very common in today's technological landscape. Capturing in-the-air body movements is more complex but can also be achieved with capacitive sensing. Although having its origins in the early 20th century [Gli00], capacitive proximity sensing has not yet been widely adopted for sensing body movements. There are numerous application examples for employing capacitive touch sensing, while proximity sensing currently evolves from a niche. Compared to capacitive *touch* sensing, capacitive *proximity* sensing requires a significantly better resolution. Therefore, proximity sensing can be regarded as the 'big brother' of touch sensing and is based on the same sensing and measurement principles. Many capacitive touch sensors are able to detect the proximity of body-parts, while being limited in terms of resolution and range.

Due to the limited interaction distance of less than half a meter, it is very hard to capture a human's whole body configuration. This would require sensors being deployed around the human, which is not feasible in many application scenarios. Considering intelligent furniture, placements near all body parts can be realized, which allows for realizing smart couches that capture the whole body configuration. In other cases, sensors are usually deployed under surfaces to detect the proximity to a body part. So far, work in this domain has been mainly focusing on recognizing foot [BHW12], hand and arm movements [Smi96].

The first applications to sense body movements with capacitive sensors in Human-Computer Interaction date back to the mid-90s. Approaches have been made at MIT to determine the position of a human hand with capacitive proximity sensors [ZSP*95]. In this paper, the authors present the "Fish" evaluation board, a board operating with shunt-mode proximity sensing. One transmit electrode and four receiver electrodes can be connected to the board. The authors briefly describe possible application scenarios of a large bandwidth. For example, two use-cases for smart furniture are presented that comprise a smart office chair for posture recognition and a smart table which can detect gestures. The presented use-cases had a reasonable follow-up at MIT with works by Smith et al., who presented more versatile architectures for hand-position recognition, for example the "'School of Fish"' [Smi96, SGB99]. This architecture is a flexible array of multiple sensors that can conduct shunt-mode measurements. In order to demonstrate the applicability of these approaches, exemplary applications have been developed that allow for manipulating objects in 3D space. In my dissertation, and also in my master's thesis, such works inspired me a lot, and many electronic circuits are derived from these works in the mid-90s.

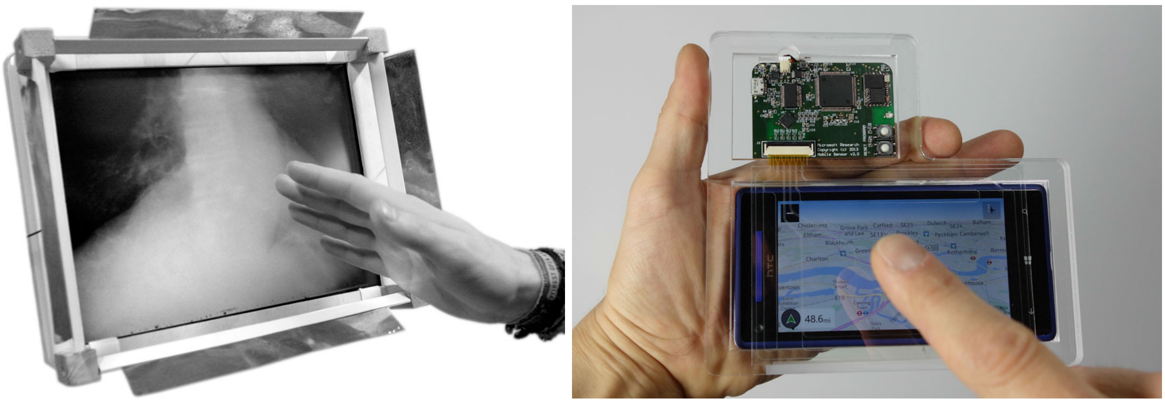

Figure 2.15.: Two approaches for recognizing body movements in front of displays from 2006 [WHKS06] (left) and 2014 [LGTIK14] (right).
[WHKS06] © 2006 IEEE.
[LGTIK14] © 2014 Association for Computing Machinery, Inc. Reprinted by permission.

In 2007, Wimmer et al. presented CapToolkit, a rapid prototyping toolkit for capacitive proximity sensing [WKBS07a]. The toolkit supports up to eight sensors that can be connected to a controller board with a USB connector. In contrast to the "Fish"-architectures presented by MIT, the toolkit does not use shunt-mode, but loading-mode sensing instead. This allows for more distributed sensing setups, in which the electrodes are deployed more sparsely. When using shunt-mode sensing, the signal decreases with the inverse cube of distance. In loading-mode sensing though, only one electrode is used for conducting the measurement, and therefore, there is no necessity to place two or more electrodes closely aside each other. Based on CapToolkit, the authors developed a number of different use-cases comprising interactive tables and an interactive shelf. The shelf is depicted in Figure 2.16

Recognizing hand movements in front of a display was presented as a small case study in [ZSP*95], illustrated in Figure 2.15. The authors realized a 2D finger-pointing that can recognize pointing gestures. It utilizes two electrodes for receiving a signal that is coupled into the persons body (transmit-mode sensing). A third electrode located at the person's left thumb enables for selecting contents, similar to a mouse-click. In 2006, Wimmer et al. presented Thracker [WHKS06], which uses loading-mode instead of transmit-mode. It utilizes four electrodes arranged around the screen to sense pick-and-drop gestures and hand locations.

Recently, Le Goc et al. transferred shunt-mode sensing to a smartphone to detect gestures atop of it [LGTIK14]. The approach is shown in Figure 2.15, illustrating the use of such a device. The portability of the setup was a challenge to signal processing as the hand posture is very variable. This induced the authors to present a localization method for the fingertip location based on regression. It utilizes Random Decision Forests with a large ground-truth database for mapping the 5 sensor values to a 3D fingertip location. Other methods for object recognition apply weighted average calculations [Bra09] or forward models [Smi96].

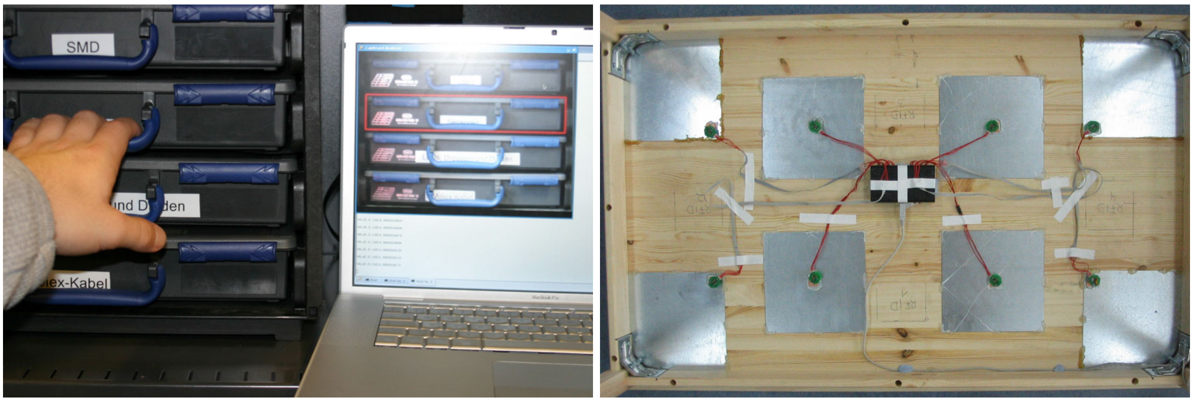

Figure 2.16.: CapShelf is a system that enables technicians track the usage of boxes with shelves [WKBS07b] (left). CapTable tracks movements atop of a table, supporting up to two persons [WKBS07a] (right).

Larger setups for recognizing body postures mainly involved person localization and tracking. With [ZSP*95], the authors were the first who designed a person-sensing room based on capacitive proximity sensing. They utilize a single large transmitter electrode underneath the floor and four receiver electrodes. This makes the user to a transmitting electrode and enables for determining the user's position. Later on, this idea was modified by Valtonen et al. [VMV09] who demonstrated the applicability in a real living environment. Here, a signal is modulated individually on each floor tile which is transmitted by the user to a receiver mounted on the ceiling. This makes it possible to track multiple users within a living environment. Steinhage et al. present a capacitive localization system with self-capacitance sensors deployed underneath a floor's surface. The interactive tiles measure capacitance to persons in proximity and can be used for recognizing falls. CapFloor [BHW12] overcomes the problem of integrating such active electronic components under the floor. It comprises a simple grid of wires with sensors attached near the walls. This makes the system more fault-tolerant, as it can be repaired with far less effort. Cohn et al. inverse the concept of using the body as a transmitter, as presented previously [CMPT12]. Instead, the authors use the human's body as a receiver electrode for environmental noise. A portable device attached to the user analyzes the induced capacitively-coupled signal from power lines and enables for gesture recognition and localization within a room.

2.3.3. Object Usage

Detecting object usage via capacitive coupling has not been investigated in such depth as recognizing body movements or interactions with surfaces. Among the reasons is the inability of pure capacitive sensing approaches to recognize the identity of objects. This is due to limited resolution and the ambiguity of sensor readings that includes physical dimensions, conductivity and proximity of the object. When posing strong initial assumptions about the type of object though, it is for example possible to detect if water bottles are filled or not [WKBS07b]. Most approaches for determining object usage are based on some sort of communication between the sensor and tagged objects. For example, Zimmerman presents a method for realizing Personal Area Networks (PANs) with objects that communicate through the human body [Zim96]. Based on this idea, the authors of [PKS*06] realized data transfer between a camera and a printer when touching both objects.

Figure 2.17.: Vu et al. communicate capacitively with touch-identification tokens to a touchscreen by spoofing touch events [VG13].
[VG13] © 2013 IEEE.

While conventional touchscreens are highly adapted to sensing touches, the perception of objects based on electric fields is still an open research question. As shown in Figure 2.17 Vu et al. design personal-touch identification tokens that can communicate with a touch-screen through air, direct contact or intrabody [VG13]. Here, the tokens generate a strong electric field to simulate touch events and communicate in this manner with the touch-screen. The token can be used for direct authentication with infrastructure, for example as a parental control. The authors achieve data rates of less than 5 bit/s.

An indirect way of sensing object usage can be based on measuring electro-magnetic interference. EMI can be measured on power lines by analyzing the spectrum on a powerline up to 500 KHz. Switch-mode power supplies or CFL lights provide very characteristic signatures at specific frequencies, which allows for distinguishing devices with high accuracy. This approach has been realized by Gupta et al. [GRP10] who could classify different active devices with an accuracy of 93 %.

2.4. Competing Technologies for Environmental Perception

In this section, I will present technological alternatives to capacitive coupling techniques. The primary goal of the section is to discuss the technologies in the context of perceiving the environment, especially to capture human interactions. I will move from small interactions, for example capturing singular fingers, to whole-body interactions. As the focus of this dissertation lies on environmental perception, I continue to distinct between the sensing goals of sensing interactions on surfaces, whole-body movements and object usage in this section. Later, the technologies are compared to capacitive sensing and communication under this respect.

2.4.1. Cameras

In terms of resolution, 2D- and 3D-depth cameras outperform many other technologies which are presented in this overview. The interaction range varies from detecting fingers and fine movements with resolutions of less than 1 mm [Lea14,WBRF13] to larger setups in which the whole body is captured and interaction takes place in room-level distances [TTW*05]. Camera systems can cover different parts of the spectrum, ranging from visible light to the far-infrared region. They can also be stationary or movable, for example having adjustable intrinsics or being mounted on a movable object. Beside restoring a 3D image from multiple points of views, it is possible to use additional modalities like time-of-flight or angle-of-arrival. Using active projections, for example in the infrared frequency region, is a widespread way of obtaining depth information.

2.4.1.1. Interactions on Surfaces

A very common way to sense touches and hand movements above a surface is the method of frustrated total internal reflection (FTIR) [Han05]. In order to achieve internal reflection infrared light is send through a material, which reflects it at its boundaries. When a finger or an object taps onto the material's surface, the light is frustrated and can be caught by an infrared camera deployed underneath. Several other techniques have been proposed which achieve a similar effect, such as Diffused Illumination [MR97].

Using projections on *transparent*, non-diffusing, surfaces, it is possible to capture hand and finger movements behind the surface. This approach can be used to scan documents, or sense gestures in front of the screen [Wil04]. Recently, Holz et al. presented a screen which is able to sense fingerprints [HB13] (Figure 2.19). It uses a fiber optic plate which reflects light in a way that it sharpens the contrast of touching objects. In combination with a high-resolution camera, users can be authenticated while touching the screen. An alternative method for touch sensing on *non-transparent* surfaces is based on shadow-shape analysis above a table [Wil05]. Here, the camera is placed above the sensing surface and detects a decreasing shadow area at the finger when a hand approaches a surface. The goal of visually recognizing objects above surfaces resulted in the development of Microsoft PixelSense technology [Mic14b], which acts in the near-infrared region.

2.4.1.2. Body Movements

Moving away from sensing touches related to surfaces, cameras are especially suitable for recognizing in-the-air gestures and finger movements. Wear-Ur-World [MMC09] presents an approach based a body-mounted camera and colour markers attached to a person's fingers (Figure 2.19). In combination with a wearable projector, persons are able to control projected user interfaces on any kind of surface. Very recently, the Leap Motion controller [Lea14] was released as a commercial product, providing sophisticated hand tracking in distances up to 1 m. In standard setups, the device is placed on a table and projects a grid of infrared points to retrieve depth information. Considering dynamically moving objects it provides an average accuracy of 1.2 mm and less than 0.2 mm for static objects [WBRF13].

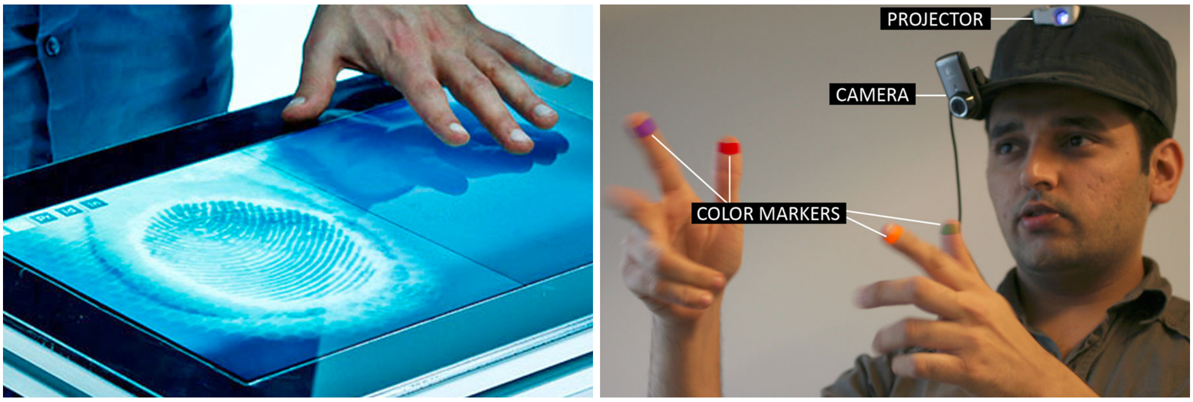

Figure 2.18.: The reflective properties of different surfaces allow for recognizing detailed information about touching fingers, such as in Fiberio [HB13]. Portable in-the-air gesture-recognition can be realized using marker-based finger position tracking [MMC09].
[HB13] © 2013 Association for Computing Machinery, Inc. Reprinted by permission.
[MMC09] Reprinted by permission of the authors.

Capturing whole-body movements with cameras can be based on a wide variety of sensing approaches, including 2D or 3D camera systems. The authors of [TTW*05] use a single RGB camera to detect falls by people using an feature extractor that recognizes the aspect-ratio of a human body. This allows for detecting different states ranging from standing, walking to falling. Very similar to the previous approach, Rougier et al. apply an ellipse for recognizing the body posture [RMSAR07]. Movable camera systems can extend the range of people localization and tracking, having a positive effect on fall recognition [LP09].

Modern 3D-cameras like the Microsoft Kinect [Pre10] currently represents a quasi-standard in whole-body tracking. The Kinect uses a simple RGB camera in combination with a projected infrared grid to derive depth information. Merging the data acquired by the grid and the RGB data enables to conduct fine-grained pose estimation for multiple persons. Since its first availability, many applications have been realized based on the two cameras that include human pose estimation [MSB11] and gaming [BSB11].

Figure 2.19.: Using large scale surfaces that generate reflections when being touched enable Braentzel et al. to derive body postures and identify users [BHH*13].
[BHH*13] © 2013 Association for Computing Machinery, Inc. Reprinted by permission.

Although frustrated total internal reflection (FTIR) is mostly used for multi-touch sensing, it has also been extended to the world of whole-body interaction. Braentzel et al. acquire a camera image from a large FTIR-enabled floor surface by placing a camera underneath it [BHH*13]. Using inverse kinematics enables them to reconstruct user poses and identify users by their footprint, as depicted in Figure 2.19. The installation intends to demonstrate the great potentials of a floor with high resolution pressure sensing. In real life though it will not be feasible due to the huge spatial constraints.

2.4.1.3. Object Usage

Recognizing objects with cameras is very similar to recognizing hands or human body parts. Therefore, the ability to detect objects in addition to human body parts does not require additional sensors but comes with the drawback of a higher computational effort. Especially when using surface-based interaction approaches [Han05], objects can be recognized very efficiently using marker-based methods. For example, Reactable [Jor10] recognizes objects based on fiducial markers [KB07] to create interactive music experiences. The manipulation of objects, in this case the position and rotation, can be captured to enrich the interaction capabilities. Moving from screens to a simple tabletop surface, objects with markers were used as PC-application controllers [CLC*10]. In larger setups, fiducial markers are commonly used in augmented reality applications, for example in combination with a mobile phone [Fia05].

As markers introduce a certain effort for tagging objects and may make use of those impracticable, purely appearance-based methods can be regarded as more realistic in the future [MGS08]. Using this approach, Molyneaux et al. identify smart objects visually to project a tangible user interface onto their surface [MG09]. In implicit interaction, the authors of [LRF12] employ a Kinect-style camera [Pre10] to recognize objects also purely on their appearance. This enables the authors to classify fine-grained activities, such as different actions while cooking.

2.4.2. Acoustic Sensing

In the last decade, acoustic sensing has proven to be a very unobtrusive technology for sensing interactions within a human's environment. In contrast to ultrasound, acoustic sensing does not rely on actively transmitted signals. Although speech recognition is a well established research field, acoustic sensing has not yet pervaded other areas on a wide scale. The sensing method is based on the fact that almost every interaction with the environment leaves acoustic traces. This ranges from tapping onto a button, moving along a surface, or interacting with devices such as water taps.

Solid materials are very well-suited for transporting acoustic noise along their surface, making them an ideal candidate for passive acoustic sensing of humans and their environment. Choosing a suitable technology for sound-pickups is not straight-forward, developers face challenges in frequency responses and coupling to materials that strongly depend on the use case. Techniques range from using microphones, microphones combined with stethoscopes, piezo-electric vibration elements, to accelerometers (MEMS) [HTM10].

2.4.2.1. Interactions on Surfaces

Sensing interactions with a wide variety of materials has been applied for sensing hand movements and touches. Even when using walls as a material for conducting sounds, ranges up to 8 m from touching the surface to the microphone can be achieved [HH08]. In order to achieve these reasonable high resolutions, a common technique is to use stethoscopes with an attached microphone. When scratching along a surface with a finger nail, sound frequencies up to 3 KHz are generated and conducted within the material [HH08]. Using this approach, finger and hand movements along walls, tables, screens, and even fabric can be recognized.

Detecting knocks, scratches and taps can be implemented very reliably using acoustic sensing [PLCH02]. Attaching four microphones to the corners of a screen, having a length of 1 m, enables to detect the point of a tapping finger with a resolution of 2-4 cm. Measuring the time-of-arrival can act as a reliable method for determining the finger position. When capturing interactions with devices, such as a screen, a lot of information is lost when touches are just analyzed by means of binary decisions [LJJ11, PH14]. Using acoustic methods, different force levels of finger taps can be distinguished [PH14]. When distinguishing two different force levels, the authors reached accuracies ranging from 99 % for two force intervals to 58 % for six force intervals.

The human body is also a well-suited medium for conducting sound. The authors of [HTM10] use an array of piezo-electric vibration elements to an upper arm. By measuring the acoustic waves propagating on the human skin it is possible to recognize different tap locations on the forearm. The resulting signals differ in amplitude and frequency that enable a user-dependent classification. Here, picking up sound with frequencies of up to 100 Hz was utilized to infer up to 10 discrete positions with an accuracy of 81.5 % on the user's skin. Reducing the number to only four locations increases the accuracy to 91.2 %.

2.4.2.2. Body Movements

In order to detect whole-body movements, one can distinguish between body-mounted and stationary acoustic sensors. In the domain of body-mounted sensors, multiple physical parameters can be captured using sound. For example, using the sound produced by muscle fibers enables Yamakawa et al. to detect different finger movements [YN12]. In contrast to muscles, conducting sounds through bone also allows for recognizing bending of a human elbow [TITO11]. Chewing, eating and drinking activities can be captured with microphones mounted inside a human ear [AJT05] or a stethoscope placed near the throat [YT12].

Popescu et al. employ acoustic sensing to detect fall situations [PLSR08]. Based on multiple microphones, the height of the sound-source is estimated. If the height is below a threshold of approximately 60 cm, an emergency situation can be derived. However, the authors state that the number of false positives is still too high for real-world deployments. Sound-source estimation can also be used for person localization by microphone arrays [CATC14]. Such systems provide an accuracy of approximately 1 m, while being able to detect multiple sources of sound.

2.4.2.3. Object Usage

One of the first approaches to sense objects with acoustic sensing is the intelligent ping pong table at MIT, developed in 1999 [IWO*99, XBY*11]. The impact of a ball on a table is captured by four microphones and the impact position is calculated. An enhanced game experience is achieved by interactive visualizations that are projected onto the table's surface. This method approaches its limits when multiple objects shall be recognized. SurfaceLink connects multiple devices that are placed on the same surface by gestures [GLIA*14]. For instance, carrying out a swipe gesture on a table triggers a data transfer from one device to the other.

StickEar [YNR13] aims for equipping objects and parts of the environment with acoustic sensing. The authors developed small stickable wireless sensor nodes. Attached to objects, they can passively sense interactions, e.g. with microwaves or books. Deployed on the wall, they can act very similar as Scratch Input [HH08], presented previously. The sensed data is then augmented with data from an accelerometer and a touch button. PANDAA picks up ambient sounds and uses them for ranging and localization among multiple devices. Its iterative approach enables to achieve accuracies for localization of approximately 17 cm.

2.4.3. Ultrasound Sensing

Not all interactions leave such nicely analyzable traces as in acoustic sensing. When it comes to in-the-air interactions or detecting stationary states, passively picking up acoustic noise is not sufficient anymore. In contrast to acoustic sensing, actively emitting non-hearable sounds can overcome the problem of recognizing stationary objects. Here, sounds are emitted and the reflected signal is measured by one or multiple microphones. In ultrasound sensing, one has to distinguish between active messaging nodes and passive backscattering techniques. The latter ones analyze backscattered signals to infer the position of body parts or the location of a person [GMPT12]. Signals from other devices received and transmitted by nodes allows for implementing an active messaging system.

2.4.3.1. Body Movements

Detecting finger and hand movements in free air is often achieved by analyzing a backscattered ultrasound signal. Here, worn nodes that use ultrasound as a messaging system can rarely be found. [KR09] apply ultrasonic waves to unobtrusively recognize one-handed gestures. The authors employ the Doppler effect which introduces frequency shifts in the backscattered signal of a moving human body part. A single transmitter and three receiver microphones are sufficient to determine a 3D movement. Due to the availability of ultrasound capabilities in consumer hardware, ultrasound approaches have also been ported to consumer smartphones and laptops. For example, SoundWave makes it possible to recognize gestures in front of the screen in ordinary laptops [GMPT12].

In order to detect whole-body movements, worn nodes as well as passively backscattered signals can be used. Multiple microphones allow for ranging and localization based on the angle or time-difference of arrival. Moreover, exploiting physical effects like the Doppler effect helps to detect changes within the environment. This enables to recognize body parts moving away or towards a sound pickup [GMPT12]. Tarzia et al. [TDDM09] use a similar technique to determine user presence near a consumer laptop. Using a measurement window of only 10 s results in an accuracy of approximately 96 %, discriminating the two classes of absence and presence. Extending the windows to 25 s almost exceeds perfect accuracy.

Attaching ultrasound nodes to a human body represents a very convenient way to identify and track a person. Techniques which use simple backscattering are not able to identify objects and thus leave a certain amount of ambiguity in their results. Randell et al. present an active ultrasonic positioning system that combines RF and ultrasound for object and person localization [RM01]. Relying on a secondary modality with faster wave propagation enables to reliably calculate the time-of-flight of the ultrasound signal [SHS01, HWL*03]. This technique has found wide applications in people localization and the corresponding tags became very small or integrated into devices [BLO*05].

2.4.3.2. Object Usage

Transmitting dedicated information about the object and its manipulations requires an active messaging system. One of the first approaches combine RF with ultrasound to achieve a reference for time-of-flight measurements [PCB00]. BeepBeep [PSZ12] solely relies on ultrasound generated by smartphones, that then supports ranging among objects. When more than three objects are involved, it is possible to determine a relative position based on the information about multiple devices. [MMOM13] follows a very similar approach: It enables users to wear a virtual uniform which emits sounds at specified frequencies. The system is intended to communicate the social state to other people. Sensing device location in a car was investigated by [YSC*11]: By classifying the mobile phone's position it is possible to differentiate between the driver using the phone and a passenger. Reynolds et al. [RMD07] introduce an ultrasound position sensing system for tangible objects above LCD-screens. The authors present interactive 'pucks' that communicate by emitting and receiving ultrasound and reconstruct their position.

2.4.4. Radio-Frequency

In radio-frequency, an electro-magnetic wave is transmitted through a medium, for example air. However, there might also be conditions when no real wave propagation through air is desired, for example by deliberately generating standing waves to sense interactions on a conductive medium. Using radio-frequency techniques to determine human interactions has gained increasing attention in the last years.

2.4.4.1. Interactions on surfaces

Time-Domain Reflectometry (TDR) has been originally designed to recognize breaches in deep-sea cables. Here, a very short broadband pulse is sent through the cable and reflected at discontinuities. The most obvious discontinuity is a cable break, but also changing environmental capacitances cause measurable reflections. Wimmer et al. presented an approach to measure these reflections in a wide variety of materials, as depicted in Figure 2.20 [WB11]. The authors were limited by the available hardware - they applied an old analog reflectometer and a camera to capture its output. Possible application scenarios range from floor systems for indoor localization to multi-touch sensors [HHL09, WB11].

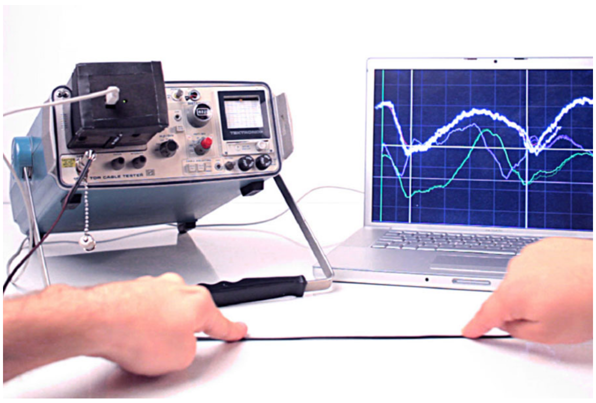
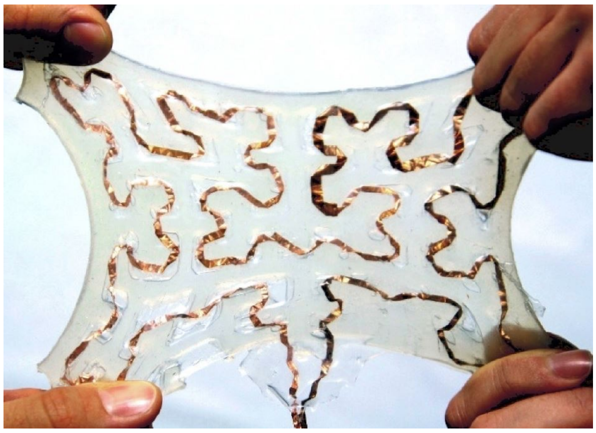

Figure 2.20.: Wimmer et al. [WB11] apply an old reflectometer for continuously recognizing multiple touches (left). Possible electrode designs include quasi space-filling Hilbert curves for multi-touch sensing based on TDR (right).

In order to sense multiple touches on two-dimensional surfaces, [DKA14] extend the concept of space-filling Hilbert curves presented in [WB11]. The primary motivation of this work was to get rid of a grid-based sensing, which is widely applied in current capacitive multi-touch displays. Although the presented concept is promising, it is an open challenge to integrate TDR into a small embedded hardware platform. The authors used an Agilent oscilloscope with TDR features.

A recent contribution by Hughes et al. [HPC14] demonstrates the use of simple hardware components to continuously locate a single finger on an antenna of fixed size. In contrast to TDR, the measurement is conducted in the frequency domain. The authors generate a standing wave and compare the phase and amplitude of the original and reflected signal. The hardware operates at a single fixed frequency of 900 MHz with an antenna of 8.25 cm length (which corresponds to $\lambda/4$).

2.4.4.2. Body Movements

Using existing infrastructure like WiFi in combination with smartphones is a common approach to indoor localization. For example, by measuring the received signal strength (RSSI), it is possible to determine an RSSI footprint of a building for indoor localization. Unfortunately, building up RSSI maps depends on the device and induces initial effort. [CPIP10] present an unsupervised approach to indoor localization, which includes clues on wave propagation and the measurements of other devices.

Applying radar with an active transmitter and receiver has been investigated by Paradiso et al. to recognize whole-body movements [PAHR97]. In this work, radar was joined with a piezoelectric smart floor to capture fine-grained foot movements. Other contributions show the applicability of MIMO radar systems which can perform through-wall imaging. This approach is especially suitable for activity recognition and localization on building level as it is not restricted to a certain room [RCP10]. WiTrack is able to recognize gestures and localize persons in indoor environments [AKKM14]. The system comprises a transmitter and three receiver antennas to leverage directional capabilities. Using frequency-modulated-carrier-waves (FMCW), the system maps shifts in time of arrival to frequency changes, which are easier to capture. WiTrack can be used for through-wall indoor localization and gesture recognition.

Recognizing whole-body movements without actively transmitting signals exploit existing radio-frequency signals. With WiSee, Pu et al. enable whole-home gesture recognition by measuring the Doppler shift generated by a moving object, like a hand or a person [PGGP13]. The authors use existing WiFi signals at 2.4 GHz having a bandwidth of 20 MHz. A great challenge is the very small Doppler shift: Movements of 0.5 m/s only resulted in changes of frequency of about 17 Hz. In order to leverage directional capabilities and locate users, the authors employed a MIMO receiver array. WiSee is able to recognize different gestures at an accuracy of 94 % without a line of sight. A conceptual drawback of Doppler-shift systems is the inability to identify static objects, like slowly moving persons [AKKM14].

2.4.4.3. Object Usage

Measuring object usage with Radio-Frequency can be based on analyzing existing wireless signals. For example, Zhao et al. use ordinary GSM signals to recognize different gestures atop of a phone [ZCA*14]. The system uses an array of receiver antennas underneath a mobile phone and measure the reflections induced by a human hand hovering over the phone. The results are very promising, the authors achieve accuracies of 87.2 % for 14 different gestures. Besides being placed in a user's palm, the system can also work in pockets, for example to put the phone to silent mode using proxemic gestures. Similar works in this group leverage gesture control abilities on everyday objects analyzing the signal strength of an RFID reader [KTG14]. The presented concept does not rely on a particular technology, any kind of existing signal can be used. The authors state that the approach can be used for gesture recognition on everyday objects and present a versatile architecture.

2.4.5. Infrared

The usage of infrared light for sensing has been partly discussed in the subsection about cameras. Infrared sensing is a widespread modality for capturing interactions with the environment. All objects that radiate heat energy generate light in the infrared spectrum. An infrared camera can be regarded as a high resolution array of infrared sensors placed in a plane. In this subsection, I will mainly discuss the usage of infrared sensors that are not placed within a plane. I distinguish between active and passive infrared (PIR) sensing. While passive sensing relies on simple sensing, active methods also radiate infrared light.

2.4.5.1. Interactions on Surfaces

Interacting on transluminiscent surfaces has been widely discussed in the section about cameras. In this domain, techniques that use infrared light like frustrated total internal reflection or diffused illumination apply [Han05]. An infrared camera mounted underneath the surface captures interactions induced by touching objects. Infrared sensors can be integrated in screen pixels [Mic14b]. They originate from coarser grids that provide vision to thin screens [IHB*07].

Considering skin as a surface, infrared sensing can be used efficiently to measure its properties or determine the distance to objects. Ogata et al. present iRing, a system based on infrared reflection sensor [OSOI12]. The sensor placed inside the ring can detect finger bending, rotation, and external force. This approach leverages gesture recognition based on finger movements.

Applying this method on a larger scale, Fukui et al. integrate an array of infrared sensors in a wristband to extract different hand contours. The reflection sensors are located side-by-side with a distance of 2.5 mm. The array provides a coarse-resolution 2D image of the human wrist contour. Fukui et al. were able to classify eight hand shapes with an accuracy of 70 %. Besides pointing on the human body, infrared sensors can also point towards moving body parts. Nakatsuma et al. apply an infrared reflection sensor to derive the position of a touching finger on the back of a hand [NSM*11].

2.4.5.2. Body Movements

In 1992, Want et al. presented an infrared localization system that makes use of active infrared badges [WHFaG92]. These badges transmit a signature to receivers that are deployed at room-level. In contrast to ultrasonic transducers, the authors promote the small form factor as an advantage.

Howard et al. investigated the applicability of an infrared proximity sensor placed on a wristband underneath a watch [HH01]. The system called LightGlove emits four infrared beams that measure the proximity to the four middle fingers. The approach is intended for realizing in-the-air or surface-based keyboards. A similar approach is followed by [LLS11], who integrate IR proximity sensors into a wristwatch. The watch can recognize gestural movements in proximity, like swiping gestures.

2.4.5.3. Object Usage

A very early example of integrating various sensors into a mobile computing device is [HPSH00]. An infrared proximity sensor in combination with an infrared LED is used to detect proximity to the device. The sensor enables proximity detection in distances up to 25 cm down to a few centimeters when direct sunlight is involved. Aitenbichler et al. present an infrared positioning sensing system for objects and humans that reaches an accuracy of 16 cm within a sensing area of 100 m^2. Butler et al. deploy infrared sensors around a device, making it possible to carry out gestures in proximity [BIH08]. Applying infrared as a sensing modality may also leverage communication abilities. This enables active objects to communicate with each other, while they sense proximity to passive objects.

The relative position between various objects can be used for context-sensitive applications, for example in tangible computing. An approach to sense relative device locations with infrared is described by [KBHG05]. Each device has a 360-degrees array of infrared transmitters that can sense the angles of arrival of other transmitters. Based on this information, a graph can be built that contains the relative locations of all devices.

2.4.6. Other Technologies

2.4.6.1. Radio-Frequency Identification (RFID)

In terms of the perceptional capabilities discussed above, RFID is only able to cover a subset of those. RFID tags can be attached to any kind of non-conductive object, allowing for identifying a tag in distances up to 100 m [Lan05]. Tags can be either passive, for example powered by an inductive near-field [NFC14], or active by employing a battery [NLLP04]. In related work, body-worn readers [SFJ*05, BLvLS10] as well as stationary readers [ZS13, NLLP04] have been used.

Although the object can be identified, RFID readers in combination with passive tags come with the drawback of a low communication range. This requires the reader antenna to be placed very close to the object, for example as a bracelet featuring a WISP tag [SFJ*05]. In order to do fine-grained activity recognition, Berlin et al. use a body-worn RFID reader and an accelerometer to identify interactions with objects [BLvLS10].

Employing WISP open-source tags [SYP*08] allows Sample et al. to integrate a variety of sensors into a tag, such as accelerometers and temperature. The authors employ a technique called ID modulation which changes the tag's id depending on the sensor measurement [SFJ*05]. In combination with ultrasound sensors, WISPs have been used for object localization, demonstrating a very high localization accuracy in the region of a few centimeters [ZS13]. However, in order to power these tags, a reader with a high output power is required to achieve ranges of up to 2.2 m. Besides identifying objects and localizing them, active RFID tags can also be attached to the human body, allowing to identify them in ranges up to 50 m [NLLP04]. By using reference tags or multiple readers in the environment, the accuracy of indoor localization can be increased. Regarding the exactly opposite way, the RFID reader can also be mounted on the object to be localized and scan tags within the environment [SN11].

2.4.6.2. Inertial Measurement Units (IMUs)

Due to the nature of IMUs, their applicability in stationary sensing setups is rather limited while they unfold their whole potentials when being attached to moving objects. Very often, IMUs are based on microelectromechanical systems (MEMS) and can be realized in very small form factors. Besides moving objects, stationary deployed IMUs are often embedded in microphones [HH08]. Therefore, I would like to refer back to the corresponding technologies presented previously in acoustic sensing.

An IMU may include an accelerometer and a gyroscope. Very often, IMUs are accompanied by a magnetic field sensor, for example to realize a compass functionality. After fusing, these multiple sensing modalities enable application developers to derive relative 3D-locations and gestures in space [Inv14]. Data from accelerometers provide a high spatial resolution in the range of $+/-16g$ with accuracies down to a few milli-g. Sampling rates of 200 Hz are not uncommon [BI04], leading to large amounts of data which have to be analyzed.

In wearable setups, IMUs have been widely used to recognize activities when being attached to different parts of a human body [BI04]. Maurer et al. [MSSD06] investigated the applicability of different placement positions of IMUs. The authors employed a three-axis accelerometer attached or placed within the pocket, bag, necklace, wrist, shirt, or belt. They classified rather elementary activities like ascending and descending stairs, walking, standing. The evaluation shows that sensors placed on the wrist and within a bag provide the highest expressiveness, while sensors placed in the shirt perform worst. Other physical setups include placements at the user's thigh, shoulder, chest, hip, or forearm [BI04]. Based on this data, researcher classified basic activities [ASLT05, BI04, MSSD06] or high-level activities [HBS07]. Activity recognition on smartphones equipped with acceleration sensors became very popular in the last years [BGC09].

2.4.6.3. Resistive sensing

Resistive sensing is only applicable in cases when the conductive properties of a material are changed based on mechanical contact or deformation. Therefore, the applicability of resistive sensing system can be primarily seen in the sensing goal of interacting with surfaces. Integrated in objects, such as a shoe, singular resistive sensors were applied to measure gait using the pressure applied to the foot [MP02, PSS*11]. A combination of resistive bend and pressure sensors may increase the expressiveness of the captured data. Creating arrays of resistive sensors is possible by applying grid-like setups, for example to analyze gait [MBBN05]. The system is based on a grid of vertical and horizontal lines that are separated by foam and generate a simple binary connection when being compressed. More sophisticated methods are based on higher-resolution pressure sensing, for example self-organizing pressure-sensing tiles by [MDL*02].

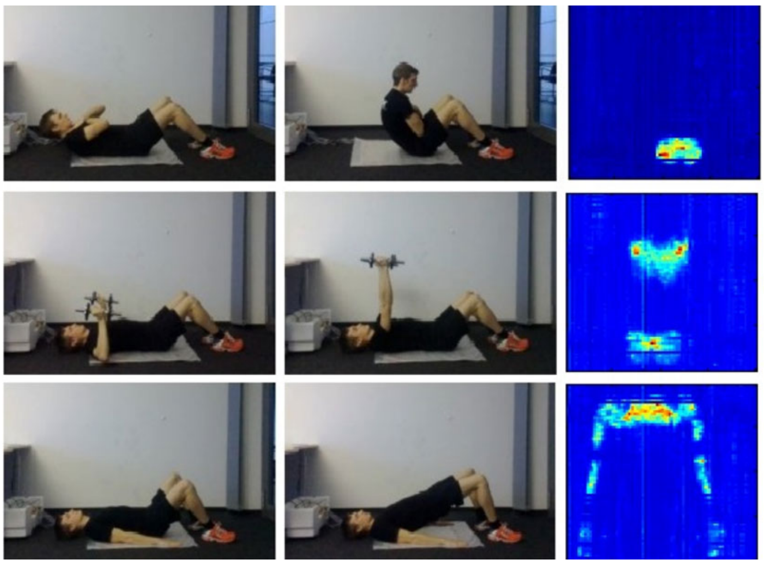

Figure 2.21.: Sundholm et al. apply a grid with higher resolution and better sensing resolution to recognize sport activities [SCZ*14].

Recent contributions in this field introduce higher spatial and temporal resolutions, for example the Smart-mat by Sundholm et al. [SCZ*14]. The authors apply a fine-grained grid with a conductive polymer between horizontal and vertical conductors. The grid lines have a spacing of 1 cm and allow for detecting activities on the mat's surface, as shown in Figure 2.21. Using a similar approach, one is able to capture interactions with objects, for example by recognizing the availability of products in retail shelves [MMFT07]. Resistive sensors can also be integrated with capacitive proximity sensors, as presented in [GSO*14]. This allows for capturing fine-grained interactions with objects using information about the touch's pressure as well as proximity to a hand.

2.4.6.4. Piezoelectric sensing

Very similar to resistive sensing, piezoelectric sensing reacts on direct mechanical influence. The piezoelectric effect introduces a voltage change when the piezoelectric material is deformed. Measuring the voltage across two points enables for detecting pressure. In microphones, piezoelectric sensing is very common, referring back to Subsection 2.4.2 about acoustic sensing [PLCH02, HTM10].

One of the first applications of piezoelectric sensing in UbiComp is Paradiso's MagicCarpet [PAHR97]. It senses human footsteps on the carpet and captures whole-body movements by introducing a second sensing modality, Doppler radar. Integrated in objects, in particular an interactive shoe, Paradiso et al. generate augmented dance experiences [PH97]. Using the piezoelectric effect is also interesting to generate energy, enabling simple push buttons to send RF messages without an additional energy source [PF01]. Paradiso's work in particular demonstrates the versatility of piezoelectric sensing, ranging from equipping large floors to small interaction surfaces in a shoe.

2.5. Discussion & Technology Comparison

In my first two research challenges, I aim at extending the perceptional capabilities of capacitive sensing. This motivation is substantiated by the unique properties of this modality. They include unobtrusive placements, low energy consumption, and computationally inexpensive data processing. Discussing and comparing properties of other technologies will make it easier to understand the limitations and prospects of capacitive sensing. Thus, taking a look aside on other modalities reveals individual strengths and weaknesses that inspired my contributions to the three research challenges. As stated in the beginning of this chapter, the choice of a suitable sensing technology is a multi-dimensional problem. There is often not just one possible and best solution, instead, each solution comes with its individual drawbacks and advantages. In this section, I will disambiguate this multidimensional problem to a certain extend, by investigating the different dimensions. I will continue to orient myself on the sensing goals of sensing interactions on surfaces, body movements, and object usage.

2.5.1. Physical Properties

In the previous sections I have shown that the presented sensing goals can be achieved with a variety of modalities. The measurable physical properties range from acceleration, to signal strengths and Doppler shifts. I will start motivating the problem of selecting a suitable physical property to measure by regarding a MEMS acceleration sensor and a GPS sensor. The MEMS sensor is able to measure acceleration, a derivative property of velocity measurable by GPS. Using the accelerometer only, it is not possible to reconstruct the velocity if no initial assumptions are made. Such initial assumptions can lead to measurement errors, even allowing errors to accumulate over time. On the other hand, the GPS sensor is able to reconstruct the velocity without greater errors, but faces drawbacks in terms of applicability in indoor environments. Figure 2.22 shows an overview of measurable physical properties.

Many systems are only able to measure changes of the environment, instead of being able to measure static conditions. Considering the example of a person who falls on the floor, a passive acoustic sensing system only has a single chance to recognize this fall. In contrast, capacitive sensors or cameras may recognize the situation even after minutes, which may lead to fewer errors.

When being exposed to a highly noisy environment, even the best sensors will fail at certain points. Different sources of noise which may affect such systems are very particular for each technology. Figure 2.23 shows an overview of possible environmental noise factors for the presented sensing technologies.

Measuring Static Conditions	Measuring Changes
Capacitance [SGB99, WKBS07a]	-
Resistance [GSO*14, SCZ*14]	-
Voltage (Piezoelectric effect) [PAHR97, PH97]	-
Velocity	Acceleration [AJT05, BI04]
RF propagation [AKKM14, HPC14, WB11]	RF Doppler effect [PGGP13]
Ultrasound propagation [RMD07]	Ultrasound Doppler effect [GMPT12]
-	Acoustic sounds [PLCH02, HH08]
Light propagation [HB13, Han05, RMSAR07]	-

Figure 2.22.: Comparison of different measurable physical properties and their derivatives. The references include exemplary works which exploit the corresponding property.

Sensing Modality	Environmental Noise Source
Capacitive	Power supplies [GPHW*14], User and electronics ground-coupling [SGB99], Skin conductivity [FL96], Conductive Objects
Cameras	Direct sunlight, Changing lighting conditions (e.g. flickering), Occlusion
Acoustic	Environmental sounds [PLCH02], Movement sounds [HTM10, HBW12]
Ultrasound	Environmental sounds [GMPT12], Multipath effects [PSZ12]
Radiofrequency	Multipath effects [PGGP13], Small moving objects (e.g. fans) [AKKM14]
Infrared	Sunlight, Changing lighting conditions, Occlusion
IMU	Sensor displacement [KL08], External vibrations (e.g. driving a car)
Resistive	Skin conductivity [FL96]

Figure 2.23.: Different types of environmental noise sources that may affect a measurement.

In this consideration, I do not include any noise that is generated by electronics, e.g. thermal noise or quantization noise.

Depending on the physical property such systems measure, they will also be able to capture the undesired parts of that property. For example, when using vision-based systems, direct sunlight on the sensor will always be an issue. On the other hand, in capacitive sensing systems, electric fields produced by switching-mode power supplies may affect the capacitance measurement. Moreover, it is hard to handle for capacitive sensing systems when the properties of the electronic's grounding change. This can be induced by partly battery-powered devices that are charged by a power supply, e.g. a smartphone or a tablet PC with a capacitive touchscreen. Regarding systems based on sound waves, environmental noise can affect measurements, such as rustling paper which produces broadband noise.

One of the prospects to enhance environmental perception with capacitive sensing is to reduce the influence of noise on sensors. This can, for example, be achieved by providing redundancy. Very often in capacitive sensing, noise is introduced on a global level, rather than on individual sensors. However, digital signal processing can still lead to reasonable inferences when fusing measurements in an intelligent way. As a contribution to the second research challenge about information retrieval from capacitive sensors, I present a method for object recognition with multiple capacitive sensors that are subject to noise [GPBKK13].

2.5.2. Placement of Sensors

Especially when comparing sensor placements to each other, the individual corresponding strengths and weaknesses in the sensing technology will be revealed. In my comparison, I adapted a classification by Michahelles and Schiele [MS04], who categorized placements in the categories of (1) installed in the environment, (2) attached to the body or object, and (3) mutual collaboration. Mutual collaboration systems require at least two components, which can be placed on the human body and in the environment. A popular example of such systems is GPS, which requires a satellite and a receiver attached to an object. Less obvious examples for mutual collaboration can be visual markers in combination with cameras, although the visual marker is usually passive.

Sensor Placement → **Sensing Goal ↓**	**Installed in Environment or in Objects**	**Attached to Human**	**Mutual Collaboration**
Interactions on Surfaces	Capacitive$^{\Delta}$ [BHW12, CBL08, CCGP13, GSO*14, GPMB11, SPH12, WKBS07a] Cameras [Han05, Wil04, Wil05, HB13] Acoustic [HH08, PLCH02] Radio-Frequency [HPC14, DKA14, WB11] Resistive [GSO*14, SCZ*14, MP02, PSS*11] Piezoelectric [PAHR97, PF01, PG97]	Capacitive [CMPT12, CGL*12] Cameras Acoustic [HTM10] Infrared [HH01]	Capacitive [VG13, YCL*11, Zim96] Cameras [Han05, Jor10, KB07, Mic14b] Ultrasound [RMD07] RFID [CPL12, FWK*13]
Body Movements	Capacitive$^{\Delta}$ [Bra09, GPBB*13, SGB99, WKBS07a] Cameras [Lea14, WBRF13, Pre10, RMSAR07, TTW*05] Acoustic [PLSR08, CATC14] Ultrasound [GMPT12, KR09, TDDM09] Radio-Frequency [AKKM14, PAHR97, PGGP13, RCP10, ZCA*14]	Capacitive$^{\Delta}$ [CMPT12, CGL*12, CAL10, GPBB12] Cameras [KHI*12, Lea14] Acoustic [HTM10, TITO11, YN12, YT12] Infrared [HH01, LLS11, NSM*11, OSOI12] IMUs [AJT05, BI04, BGC09, MSSD06] Piezoelectric [PH97]	Capacitive$^{\Delta}$ [GPHW*14, Zim96] Cameras [Fia05, MMC09] Ultrasound [AM03, HWL*03, RM01, SHS01] Radio-Frequency [CPIP10] Infrared [WHFaG92] RFID [SYP*08, CPL12, FWK*13, NLLP04]
Object Usage	Capacitive [CCGP13, GSO*14] Cameras [MGS08, MG09] Acoustic [IWO*99, XBY*11] Radio-Frequency [HPC14, ZCA*14] Resistive [GSO*14, SCZ*14, MP02, PSS*11] Piezoelectric [PAHR97, PF01, PG97]	Capacitive $^{\Delta}$ [GPBB12] Cameras Acoustic [YNR13] IMUs	Capacitive$^{\Delta}$ [GPHW*14, VG13, YCL*11] Cameras [Fia05, MMC09] Ultrasound [RMD07] Infrared [KBHG05] RFID [BLvLS10, SFJ*05]

Figure 2.24.: Comparison of different sensor placements, inspired by Michahelles and Schiele [MS04]. The fields marked with background colors indicate my scientific contributions.

Figure 2.24 shows that many placements in combination with a sensing goal are well explored fields among all modalities. For example in environmental deployments, capturing interactions on surfaces and recognizing body movements has been widely investigated. In this field, my contributions refine existing techniques that mainly target at the first research challenge on new capacitive sensing methods. In particular, I raise the spatial and temporal resolution in proximity sensing [GPBB*13].

Up to now, mutual collaboration for sensing object usage and body movements has not been widely investigated in capacitive sensing. Other technologies have proven to be very suitable in this field though, for example visual markers or ultrasound beacons. That raises the question why there are only few contributions investigating the applicability of capacitive coupling in this domain. This made me curious and motivated me to work on Capacitive Near-Field Communication [GPHW*14]. As a unique property CapNFC provides novel interaction possibilities like through-body interactions. The method is able to combine sensing and communication in an intelligent way. The work contributes to the first research challenge on new capacitive sensing methods.

Attaching sensors to humans for sensing body movements is commonly applied among many technologies. In contrast to measuring muscle contractions with capacitive sensing [CAL10], using this modality to sense the outside world also offers new prospects for capacitive sensing [CMPT12,CGL*12]. As a contribution to the third research challenge on implicit and explicit interaction, I investigated capacitive sensing in wristbands to measure the proximity to body parts and objects [GPBB12]. The retrieved data can be used to enhance activity recognition by an additional modality.

The figure also reveals categories that are less densely populated. For example, achieving the sensing goal of recognizing interactions on surfaces with body-attached sensors is quite tricky. In this case, only few physical sensing opportunities apply. In capacitive sensing, Cohn et al. detect electric potentials on the human body and which enables to recognize the proximity to power lines on the surface [CMPT12]. The same applies for sensing object usage with sensors attached to a human. Here, mainly mutual collaboration approaches using technologies like RFID apply [SFJ*05, BLvLS10]. Recognizing object usage by singular sensors placed on the human body opens new perspectives for activity recognition.

2.5.3. Range, Resolution and Energy Consumption

The supported interaction range has a significant impact on the perceptional capabilities of a sensing modality. A tradeoff has to be found between the desired range, resolution and energy consumption. To a certain extent, energy consumption scales with the processing complexity of a problem. For example, when using a camera, the processing complexity and also the energy consumption will be higher than in capacitive sensing. This is tightly coupled with the resolution, as camera images contain more information than a typical array of capacitive sensors can deliver. Also, range and energy consumption must be balanced very sensitively. Figure 2.25 shows these two dimensions and the corresponding modalities discussed in the following.

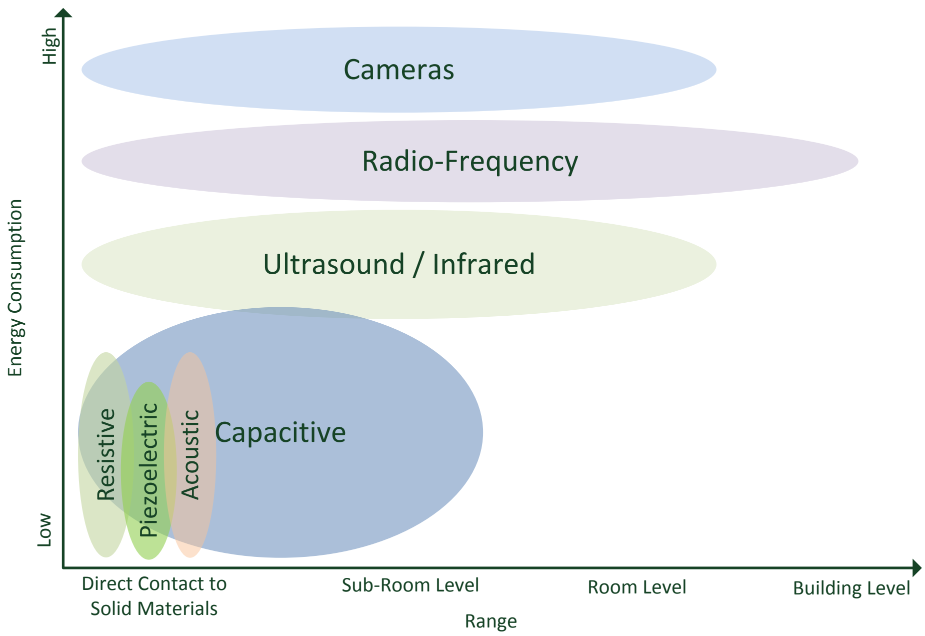

Figure 2.25.: Measurement range versus energy consumption of a single sensor including peripherals for installations within the environment.

Compared to capacitive sensing, all presented approaches that rely on propagating waves in the air (sound, light, or RF) support relatively large detection ranges. They allow for detecting approaching body parts in room-level or even building-level distances. As a drawback, the energy consumption is usually higher than in passive setups when actively emitting mechanical or electromagnetic waves. Although cameras can be either passive or active, the high resolution introduces complexity in further processing, leading to a higher energy consumption. In contrast to cameras which require a line-of-sight, RF-based sensing can support ranges on building level. A drawback of RF-based sensing is the lower resolution. Comparing passive acoustic and capacitive sensing reveals very similar properties in terms energy consumption. In acoustics, most materials conduct sound very well and microphones can thus be placed far away from the excitation source [PLCH02]. A key drawback in passive acoustic sensing is that only changes and interactions with solid materials can be detected. In contrast, capacitive sensing allows for detecting hands and fingers, even if they do not move. Employing capacitive touch sensing for determining a finger position on surfaces requires a grid of electrodes. The resolution can be very high which can be applied to distinguish multiple fingers. Easier deployments such as in acoustics have a positive impact on the solution's unobtrusiveness but provide measurements with less resolution.

This brings in the question on how to extend environmental perception with capacitive methods, taking the dimensions of resolution, range and energy consumption into account. Due to physical restrictions, it is hardly possible to extend the range of capacitive sensors to room level. Of course it is feasible to increase range by providing better electronics and using higher voltages. But when large sensing distances with a certain directivity are required, capacitive sensors will not be able to compete with technologies that rely on wave propagation. The unique properties of capacitive coupling are revealed when regarding range and energy consumption in combination. In contrast to the great bandwidth of low-power capacitive proximity sensors, the low power consumption has not been widely exploited for capacitive communications. Therefore, capacitive coupling as a very-low power, sub-room level communication method [GPHW*14] contributes to the first research challenge.

2.6. Summary

In this chapter, I discussed related work in the domain of ubiquitous interaction and environmental perception. As the focus of this dissertation is on capacitive sensing, I first introduced the exploitation of capacitive coupling in nature. Weakly electric fish have been using this modality for thousands of years to communicate with other individuals and sense their environment. I presented the first uses of capacitive sensing by humans, starting with a music instrument developed in the early 20th century. Next, I provided a conceptual background that yields for all capacitive sensing and communication setups.

In order to categorize the perceptional capabilities of sensing modalities, I introduced three sensing goals. They comprise sensing (1) interactions on surfaces, (2) body movements, and (3) object usage. Relevant sensing modalities including cameras, ultrasound, and infrared sensing, were discussed with this regard. In the following, I identified prospects to extend the perceptional capabilities of capacitive sensing. To reach this goal, I compared capacitive sensing to other technologies. For example, I analyzed different sensor placements with respect to the three sensing goals. With additional dimensions like range and energy consumption, I showed how my work contributes to the defined research challenges.

I would like to summarize the reasons that excel capacitive sensing for environmental perception. Capacitive sensors can be placed very easily on humans, on objects, or within the environment. The conductive properties of a human body can be exploited that can be used for developing highly interactive systems at low cost. Compared to ultrasound and infrared sensors, there is a tradeoff between detection range and energy consumption. On the one hand, capacitive sensing systems can achieve a lower power consumption than the latter systems. On the other hand, the range of capacitive sensors is significantly smaller. An important strength of capacitive systems is the deployment. They can be realized in small form factors with unobtrusive placements under any kind of non-conductive surface. The low power consumption enables the use of this modality in systems with low maintenance effort. In contrast to systems that achieve far better sensing resolutions, capacitive sensing requires lower computational abilities. This property also reflects in the low cost and power consumption of capacitive sensors. In conclusion, capacitive sensors do not achieve the best possible sensing resolutions and ranges, but their advantages lie in form factors, unobtrusiveness, and low power consumption. The versatility of capacitive technologies matches very well with the high complexity that is introduced in interaction design.

3. Improving Capacitive Proximity Sensing with OpenCapSense

In their very reflective article on UbiComp, Greenberg et al. identified proxemic interactions as one of the key concepts of future UbiComp systems [GMB*11]. The authors begin their argumentation by citing Bill Buxton with the following words:

> *"When you walk up to your computer, does the screen saver stop and the working windows reveal themselves? Does it even know if your are there? How hard would it be to change this? Is it not ironic that, in this regard, a motion-sensing light switch is 'smarter' than any of the switches in the computer, AI notwithstanding?"*
>
> — Bill Buxton, 1996, Living in Augmented Reality [Bux96]

Buxton made this statement in 1996, and almost twenty years after I believe that few computers behave in this reasonable and intelligent manner. When we regard new device categories like tablet PCs or smartphones, proxemic interactions slowly evolve to new intelligent interaction concepts. For example, smartphones are able to unlock themselves based on identifying the user's face in front of a screen. Taking another step towards Weiser's vision [Wei99] requires not only singular proximity-aware devices, but an environment that is capable of sensing proxemic interactions.

In the previous chapter, I have introduced various physical sensing opportunities for proxemic interactions. Many of these approaches cannot be applied off-the-shelf, they rather require knowledge in two disciplines - electrical engineering and computer science. Successful prototyping platforms like Arduino[1] or Gadgeteer [VSH11] show the potential benefits of bridging that gap.

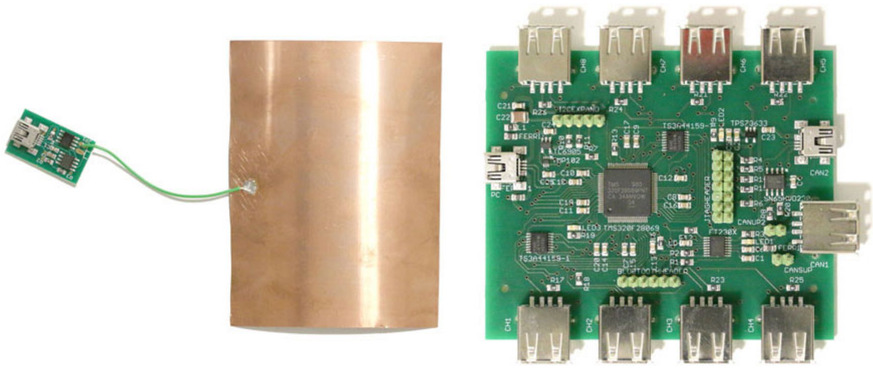

Figure 3.1.: The OpenCapSense board supports eight capacitive touch and proximity sensors. Moreover, it integrates various communication interfaces such as Controller-Area-Network for creating sensing arrays.

[1] http://www.arduino.cc/ (Date accessed: 2014/08/31)

The limitations in current capacitive prototyping platforms lead me to the development of a novel prototyping toolkit, shown in Figure 3.1. The kit, called *OpenCapSense*, enables researchers who are not familiar with electronics to quickly realize highly innovative user interfaces. It is the first toolkit which supports all three operating modes in capacitive proximity sensing introduced in Section 2.3.2. In contrast to previous work [WKBS07a] a significantly higher temporal and spatial resolution is achieved. Based on its sophisticated processing capabilities, OpenCapSense achieves measurement update rates of more than 1 kHz. The typical spatial resolution varies between one millimeter at close object proximity and around one centimeter at distances of 35 cm or above.

Figure 3.2 depicts the different dimensions of proxemic interactions [GMB*11]. Based on the contribution presented in this chapter, I am able to cover the dimensions of proximity, orientation, movement, and location. In order to obtain adequate continuous measures for these dimensions, processing algorithms are required that will be presented in Chapter 5. Due to its nature, capacitive sensing only has limited potentials to identify persons or objects based on proximity. Moreover, caused by the low resolution, sensing orientation is only possible to a certain extent. In the following chapter, I will introduce a concept to complete the picture and support all dimensions.

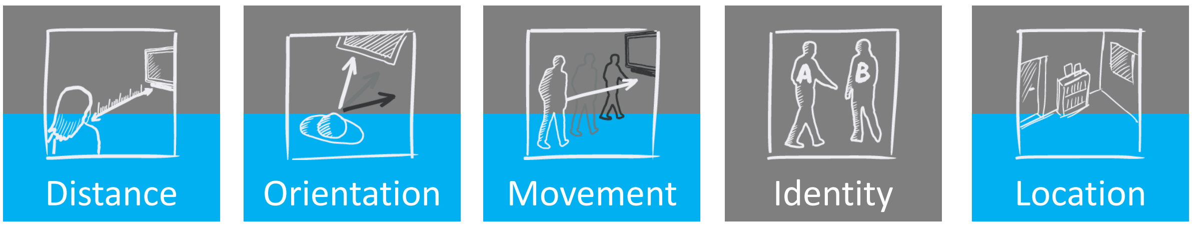

Figure 3.2.: OpenCapSense covers four dimensions of proxemic interactions introduced in [GMB*11].
[GMB*11] © 2011 Association for Computing Machinery, Inc. Reprinted by permission

This chapter is an extended version of the paper [GPBB*13] which includes evaluation results from Yannick Berghoefer's master's thesis [Ber12]. The use of 'we' in this chapter refers to the the paper's authors Tobias Grosse-Puppendahl, Yannick Berghoefer, Andreas Braun, Raphael Wimmer, and Arjan Kuijper. In this chapter, we introduce the following scientific contributions:

- We introduce a novel open-source[2] toolkit for capacitive sensing, providing eight sensor channels and two combinable sensor types. OpenCapSense covers all three capacitive proximity sensing modes: loading mode, shunt mode, and transmit mode.
- The system is evaluated with various sensing electrode sizes as well as opaque and transparent materials. It shows significant improvements in spatial and temporal resolution, compared to [WKBS07a].
- We demonstrate practical uses of OpenCapSense in different applications scenarios. These include use-cases from the area of explicit interactions, for example gesture recognition. In implicit interaction, OpenCapSense can be employed to prototype fall detection systems.

3.1. OpenCapSense Rapid Prototyping Toolkit

Especially in the field of capacitive proximity sensing, we experienced a lack in available prototyping solutions. In contrast to some electronics toolkits that support capacitive *touch* sensing [GF01, SZH12], capacitive *proximity* sensing requires a significantly better resolution. A few years ago, the CapToolKit

[2]Hardware licensed under Creative Commons CC BY-NC-SA 3.0, software licensed under GPL v3

[WKBS07a] was released, an easy-to-use prototyping platform for proximity sensing applications. It has been used by several research groups to realize applications, such as activity recognition and whole-body interaction. The kit consists of a circuit board that incorporates a microcontroller and interfaces to connect up to eight sensors that are attached using regular USB connectors. Although several tutorials on capacitive proximity sensing can be found online [3], these generally require soldering and programming.

Only few solutions with many limitations were available for prototyping capacitive sensing applications. There was no prototyping toolkit which combines all three capacitive proximity operating modes (loading mode, shunt mode, and transmit mode) in a single platform. Moreover, the spatial and temporal solution of existing toolkits was too low for use-cases like recognizing fast whole-body interactions or gestures. Similar requirements have driven the development of our predecessor, the CapToolKit, which result in analogous design decisions [WKBS07a].

We put our contribution under the hood of open innovation: Open-source hardware and software allows others to improve and customize our design. Moreover, we aim for easy integration of the system using plug-and-play operation. A simple text-based protocol can be used for configuring the device and report measurement data. The architecture should also allow for flexible and customizable electrode layouts In order to achieve the best possible performance, we aim at reducing stray capacitance and external environmental noise. This is achieved by putting the sensor circuit as close to the electrode as possible und applying shielding techniques. The system is intended to be affordable for the research community, it comprises affordable and easily obtainable hardware components.

We have identified three major limitations of the CapToolKit that have driven the design of the OpenCapSense toolkit:

- CapToolKit only supports loading mode sensing. OpenCapSense additionally supports transmit and shunt mode sensing. These modes decouple sending and receiving electrode and allow realizing applications with a higher detection range and a better temporal resolution due to parallel transmitter operation.
- The platform has a low update rate of approximately 50 Hz due to its slow 8-bit microcontroller and software architecture. The update rate is further reduced when multiple sensors are connected. OpenCapSense offers an update rate between 1 kHz and 250 Hz when using all sensors. This is sufficient for tracking very fast motion, for example in fall detection scenarios.
- CapToolKit is not suitable for creating larger sensing arrays, because it requires to individually connect each controller board to a host PC. As CapToolKit boards are not synchronized with each other, active sensors may negatively affect adjacent other sensors. We address this issue by providing two real-time bus interfaces that can be used for synchronization and data exchange.

OpenCapSense consists of a board employing a microcontroller that is tailored to complex real-time calculation and control tasks. The sensors can be combined freely with little configuration effort, enabling application developers to use different measurement modes or combine them to realize hybrid measurements. Moreover, the software provided with OpenCapSense provides functionality from evaluating sensor measurements up to realizing a specific user interface using machine-learning techniques.

3.1.1. OpenCapSense Board

The overall board architecture is shown in Figure 3.3. It provides eight sensor channels for capacitive proximity measurements. The board is built around a TMS320F28069 microcontroller from Texas Instruments. This microcontroller has a native floating point unit for fast digital signal processing, like

[3] e.g. http://arduino.cc/playground/Main/CapacitiveSensor

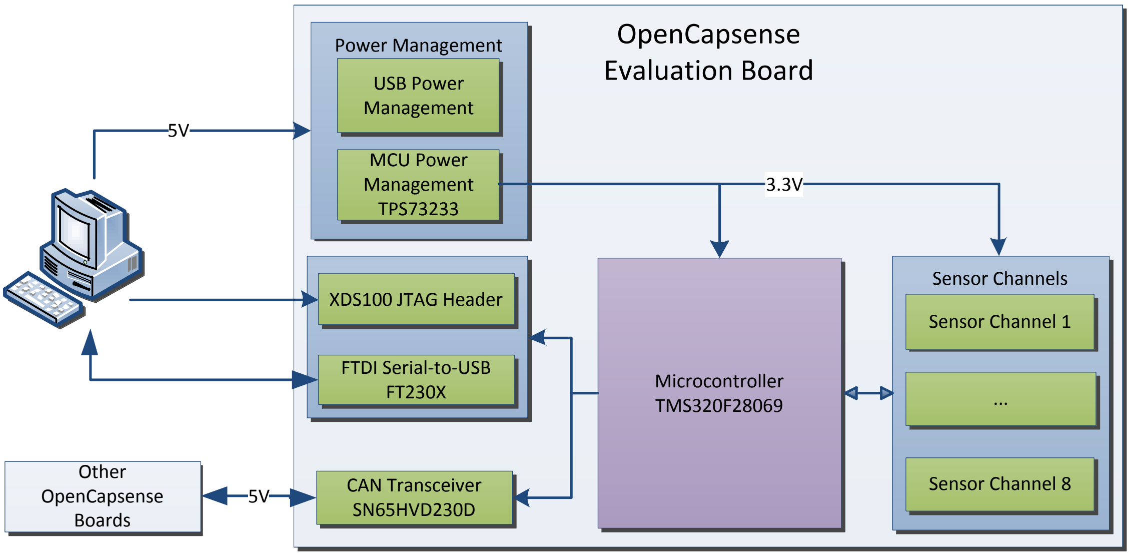

Figure 3.3.: Block diagram of OpenCapSense's board architecture. OpenCapSense employs a powerful microcontroller for real-time digital-signal-processing, such as performing FFTs. It provides eight sensing channels and sophisticated communication abilities.

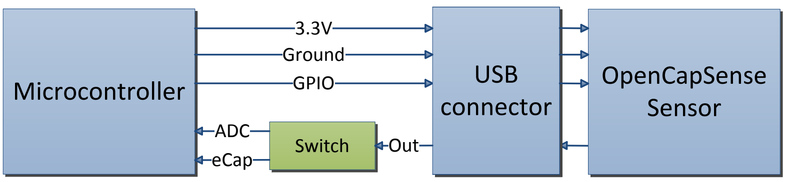

Figure 3.4.: Illustration of a single sensor channel. Each sensor has a 3.3V supply, a dedicated control pin (GPIO), and an output channel that can be connected to the microcontroller's analog-to-digital converter (ADC) or time-capturing unit (eCap).

performing Fast Fourier Transforms (FFT). The board is driven by a 5V voltage supply source that is either connected by Universal Serial Bus (USB) or by Controller Area Network (CAN). A linear dropout regulator provides a 3.3V supply to drive the microcontroller.

Sensors are connected to the board by a standard USB cable. A single sensor channel, as depicted in Figure 3.4 represents a generic interface that supports different sensor types. Apart from two power supply lines, there are two software-configurable lines for communicating with a sensor. The first line leads to a microcontroller's general-purpose-input-output port (GPIO) while the second line can be switched to either an analog-to-digital converter (ADC) or a time-capturing unit. This architecture enables us to dynamically configure the toolkit's sensor channels for different sensor types. For example, digital output signals, like frequency-modulated rectangular pulses, benefit very much from time-capturing units that can measure time intervals between the rising and falling edges of a signal. The time-capturing units analyze signals with a temporal resolution of 12.5 ns. Analog receivers are usually connected to an

ADC that samples the signal at high data rates up to 100 kHz and 12 bit resolution. This method allows for obtaining a set of values representing the signal's amplitude in a certain timespan.

Ubiquitous interfaces are often driven by a variety of different sensors, such as accelerometers or temperature sensors. Thus, it was a key requirement for us to support the integration of other types of sensors. These sensors can be interfaced using an expansion header driven by an Inter-Integrated Circuit (I^2C) bus. The board also includes a Controller Area Network (CAN) interface for realizing larger sensing arrays and simplify prototyping in automotive applications, where CAN is ubiquitously used. The board can be connected to a computer using USB, providing both power and a communication interface. As soon as the board is connected to a computer, it will open a virtual serial port and start sending sensor data. This data can be easily processed using third-party libraries or our own programming interface. It is also possible to send control information from the computer to the board, for example to configure sensors or perform a system reset.

3.1.2. OpenCapSense Sensors

We have created two exemplary sensors that can be applied to realize all measurement types. Loading mode sensors are especially suitable for implementing larger sensing systems for activity recognition or whole-body interaction [WKBS07b]. Shunt mode sensors can be applied to realize gesture recognition interfaces with a high spatial resolution [GPMB11]. Using transmit-mode measurements, we can distinguish between different users. This measurement mode can be realized applying our shunt mode sensors.

3.1.2.1. Loading Mode Sensor

The loading mode sensor is based on a timer configuration called astable multivibration [Gat06]. The timer controls the charging and discharging cycles of the capacitor that is created by the sensing electrode and the surrounding environment, e.g. a person's limb or body. It toggles succeeding charging and discharging cycles at the time when an upper or lower threshold voltage at the virtual capacitor is reached. In order to measure the capacitance in a certain direction and prevent disturbances from objects nearby, a shield electrode can be placed directly underneath the measuring electrode. The shield is driven with the same potential as the sensing electrode, such that the capacitance between the two electrodes is negligible.

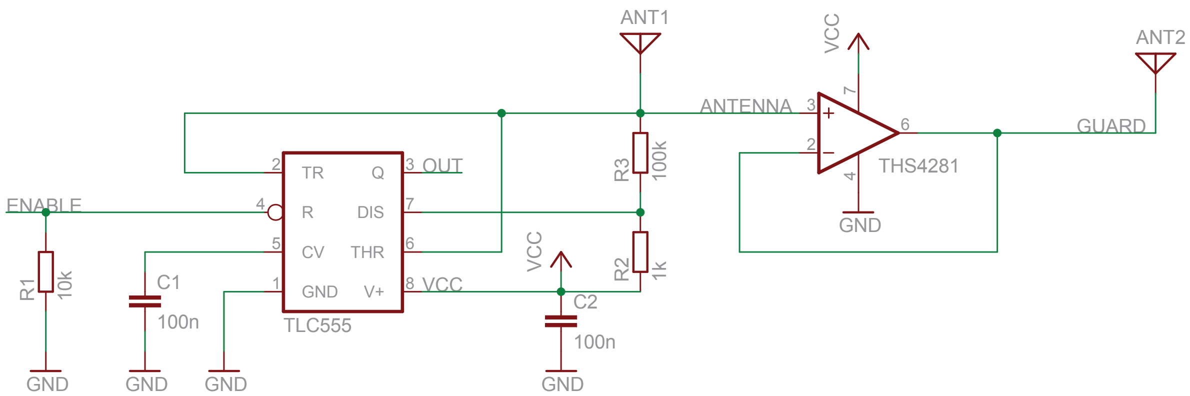

Figure 3.5.: Schematics of the astable multivibrator sensor.

The schematic of a OpenCapSense's astable multivibrator sensor is shown in Figure 3.5. It utilizes a standard 555 timer (TI LMC555) which controls between charging and discharging the capacitor C, while its size depends on the proximity of an object and the setup's parasitic capacitance. Here, the total period is defined by the RC constant, which was introduced in chapter 2.2.3.4. The oscillation frequency f is given by [Tex14a]:

$$f = \frac{1}{T} = \frac{1.44}{(R_2 + 2R_3) \cdot C} \Leftrightarrow C = \frac{1.44T}{(R_2 + 2R_3)}$$

Figure 3.6 illustrates the timer's operation. The capacitor is slowly charged with the current flowing through R_3 and uncharged with the current flowing through $R_{2,3}$. A trigger switches between charging and uncharging when 1/3 or 2/3 of the maximum voltage is reached. Thus, when the capacitor's value increases, charging and discharging cycles take longer. The measurement time is then converted to a digital representation at the corresponding microcontroller.

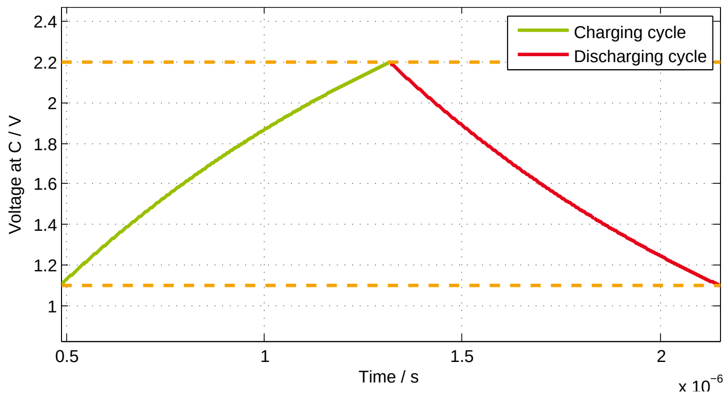

Figure 3.6.: The timer triggers succeeding charging and discharging cycles, that discharge the capacitor to 1/3 and 2/3 of maximum voltage. This curve was simulated with $R_3 = 100k\Omega$ and $C = 12pF$.

Based on that, we calculate the maximum achievable resolution in one measurement cycle. Therefore we compare two cycles of length T_c and $T_c - T_{res}$, where $T_{res} = 12.5ns$ is the temporal resolution of OpenCapSense's microcontroller. We compare the resulting two capacitances to retrieve our resolution ΔC:

$$\Delta C = \frac{1.44}{(R_2 + 2R_3)} \cdot (T_c - (T_c - T_{res})) \approx 89.5fF$$

Referring back to the plate capacitor model, given in Equation 2.6, this would correspond to detecting distance changes of $\Delta d \approx 0.01m$ at distance $D = 0.1m$ with $A = 0.01m$:

$$C = \varepsilon_0 \varepsilon_r \frac{A}{D} \text{ and } (C - \Delta C) = \varepsilon_0 \varepsilon_r \frac{A}{D + \Delta d}$$

$$\Delta d = \varepsilon_0 \varepsilon_r \frac{A}{\varepsilon_0 \varepsilon_r \frac{A}{D} - \Delta C} - D$$

In order to increase this coarse resolution, multiple cycles are performed sequentially for each measurement. Moreover, in this simple consideration measurement noise is not included. In the next section, experimental data will be presented and compared.

3.1.2.2. Shunt Mode Sensor

In shunt mode, a receiver electrode is used to measure the displacement current from a transmitter electrode [SGB99]. When a human body part enters the electric field, the field between a transmitter and receiver is interrupted. This results in a decrease of displacement current and thus a decreasing capacitance between the transmitter and receiver. Due to separate transmitter and receiver electrodes, shunt mode offers the possibility to perform parallel measurements using multiplexing approaches [SGB99]. The shunt mode sensor measures the displacement current floating from a transmitter electrode to the receiver electrode.

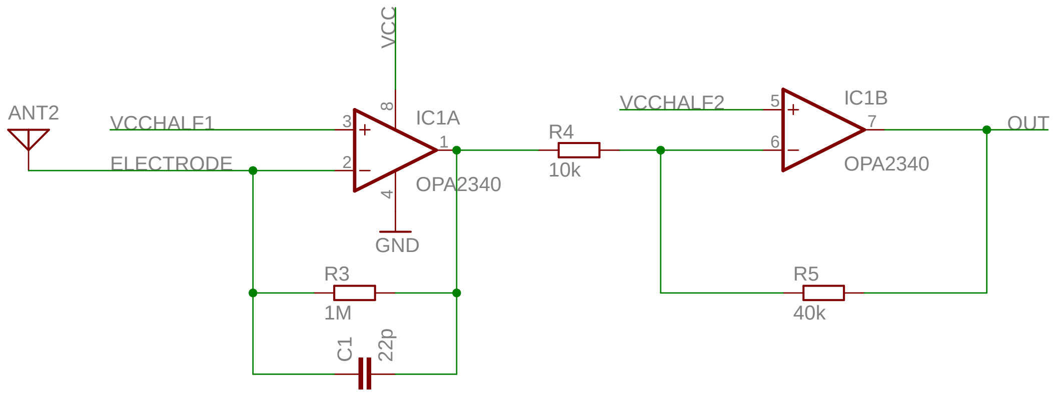

Figure 3.7.: The shunt mode sensor integrates two operational amplifers for receiving the displacement current and for active filtering.

Figure 3.7 depicts the hardware setup for the shunt mode sensor, employing two operational amplifiers (TI OPA2340). The setup's first stage is responsible for amplifying the incoming displacement current with additional low-pass filtering achieved by adding $C1$. The configuration is known as a transimpedance amplifier. Since no symmetric power supply is available, we use half of the supply voltage as an analog ground [SGB99]. $R3$ is used to control the gain factor of the first amplification stage. In the following, a gain stage amplifies the signal by a factor $G = R5/R4$. The signal can then be quantized by OpenCapSense's microcontroller. Due to the different behavior of hardware components in terms of transimpedance gain, the sensor's resolution will be determined experimentally in the following.

3.2. Evaluation

In order to analyze the capabilities of the OpenCapSense toolkit, we have conducted different experiments investigating the behavior of our loading and shunt mode sensors. There are two essential factors that can characterize this behavior. The first factor is the maximum spatial distance between a sensor electrode and a detectable object at which a presence might be registered. The second factor is the signal-to-noise ratio that can be affected by electromagnetic interference, thermal noise or ambient temperature variation. We combine both factors to derive the spatial resolution of our system.

Considering the broad spectrum of applications for capacitive sensors, we also investigated different electrode materials. Copper, which combines excellent surface resistance and low cost is the material most commonly used for electrode design. However, in many applications, such as gesture recognition, it is desirable to realize capacitive proximity sensing on transparent surfaces like windows or displays.

Various materials can be applied in those scenarios, an evaluated subset is depicted in Figure 4.7. For example, a foil of polyethylene terephthalate (PET) coated with indium tin oxide (ITO) provides both

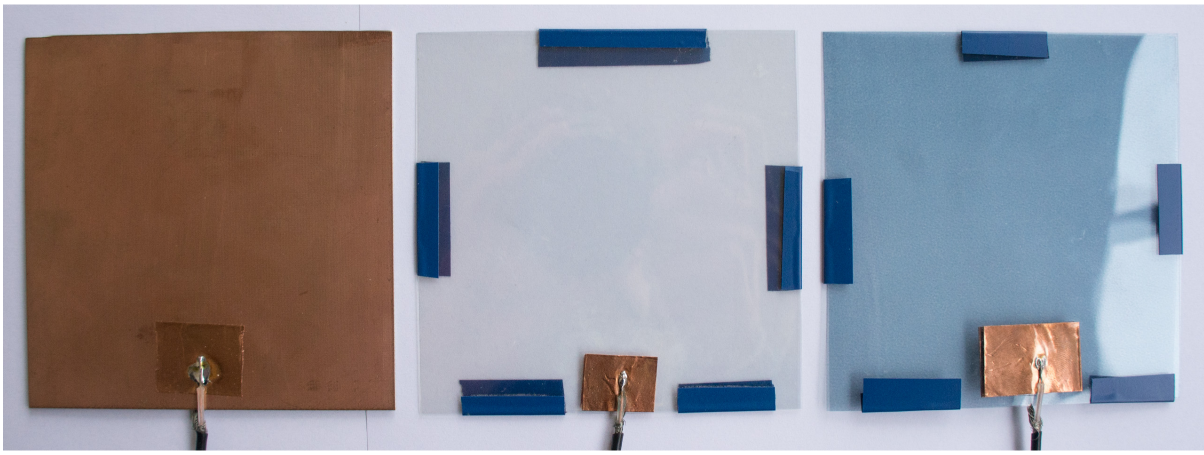

Figure 3.8.: Three different materials (copper, indium tin oxide and PEDOT:PSS) that can be applied as sensing electrodes [Ber12]

high conductivity and high transmittance of light at visible wavelengths. Another interesting material is PEDOT:PSS, an electrically conductive material made from polymers. It is used in different research disciplines like printed field-effect transistors and photovoltaic cells. Since it can be applied to different materials using ink-jet printing, it promises very good properties for rapidly prototyped sensing electrodes [CJC*06, BCS*12]. Both variants can be used to create printable transparent conduction films that offer manifold alternatives for sophisticated sensor designs and rapid-prototyping. Figure 4.7 show three different electrode materials that are investigated. ITO shows much better properties in terms of transparency, while it is easier to apply materials like PEDOT:PSS, as they can be inkjet printed.

Material	Transparency	Color distortion	Spatial resolution	Rapid-prototyping abilities
Copper	–	–	++	+
ITO	++	+	+	o
PEDOT:PSS	o	o	+	+

Figure 3.9.: Properties of transparency for different prototyping materials. ([++] very good, [+] good, [o] satisfactory, [-] unsatisfactory) [Ber12]

Figure 3.9 depicts the properties of two transparent prototyping materials [Ber12]. As a reference, copper is compared to the these materials. Although PEDOT:PSS provides a slightly lower resolution than ITO, prototyping is easier since it supports inkjet printing. In order to adapt electrode layouts with ITO, we used hydrochloric acid to etch a custom electrode structure on an ITO-covered PET sheet. Moreover, ITO is very sensible to bending and movements, which requires a very sensible handling while prototyping applications.

3.2.1. Measurements

Based on the principles introduced in section 3.1.2, we have to specify the relation between objects in a sensor's detection area and its measurement output. Regarding the loading mode measuring principle, the sensor output has its minimum value if no object is present within its detection range. In this case, the capacitance and the respective charging time of the capacitor is small. The sensor output reaches its maximum at very small object distances, when the capacitance is high.

The measurement output of a shunt mode sensor shows different characteristics than the one of a loading mode sensor. The output is proportional to the displacement current that floats into the receiver electrode. Therefore, the maximum sensor value is obtained when no object is close to both electrodes. When an object approaches and disturbs the electric field between the transmitter and receiver, the displacement current and the measurement output decrease. The loading mode and the shunt mode sensor's raw sensor values are normalized to a range between 0 and 1 to simplify the processing of measurement data of different sensor configurations.

3.2.2. Test setup

As mentioned in the introduction, the capacitive proximity sensing toolkit presented by Wimmer et al. [WKBS07a] is a commonly used platform for implementing capacitive loading mode applications. Only few capacitive sensing systems were evaluated in a quantitative manner by presenting data related to the sensor detection range and spatial resolution. In their work about CapToolKit, Wimmer et al. adopted a measuring arrangement presented by Zimmerman et al. [ZSP*95]. Considering this well-defined test setup, we evaluated OpenCapSense's and CapToolKit's capabilities under comparable prerequisites and circumstances that are described in the following.

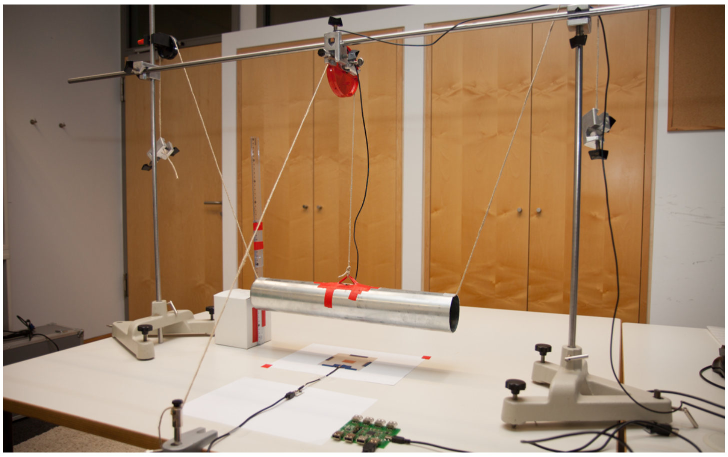

Figure 3.10.: Test setup to determine loading mode sensor's resolution and behavior related to different electrode materials and sizes. The tube acts as a surrogate arm and can be lifted up and down.

Regarding the different properties of human arms and hands, e.g. shape and conductivity, it is obvious to use a standardized measuring body. Such an object should have properties similar to a human arm and be easily applicable to different experiments. Zimmerman et al. [ZSP*95] introduced a grounded aluminum tube with a length of approximately 48 cm and a diameter of 8 cm that acts as a surrogate arm. They showed that this tube can be an appropriate replacement for a human arm, delivering slightly deviant measurement results. Based on this description, we developed a test setup that is shown in Figure 3.2.2.

The tube hangs floating freely, parallel to the ground surface, on an insulating cord, which runs over a fixed roll to adjust the tube's distance to the ground. The tube is stabilized by an adjustable second cord that runs through the tube. The tube and the measuring rack are grounded to the neutral conductor of a power socket. The sensor electrode is placed at the tube's center perpendicular to the ground.

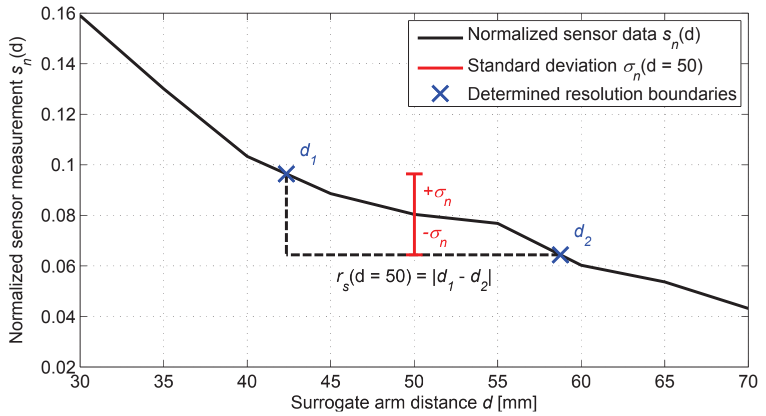

Figure 3.11.: The approach to determine the spatial resolution of a sensor configuration is based on a linearly interpolated measurement series and the standard devition $\sigma_n(d)$ at a given point.

3.2.3. Spatial resolution

As the descriptions of previous measurements [WKBS07a, SGB99] provide only little information on how the resolution of a capacitive proximity sensor can be measured, we introduce a principle to determine a system's spatial resolution. The normalized sensor values related to the distance of the aluminum tube are determined with a series of measurements. We adjust the aluminum tube to distances from 0 cm up to 40 cm. For every distance d, a series of samples is recorded for 10 seconds. This allows calculating basic statistical values of the sampled sensor measurements. The most important values are the arithmetic mean $\bar{s}_n(d)$ and the standard deviation $\sigma_n(d)$, the latter being the main criterion to determine the system's signal-to-noise ratio. The lower $\sigma_n(d)$ for a measurement series at a distance d, the higher the certainty about the presence of an object at a certain distance. Therefore, the spatial resolution strongly depends on the standard deviation.

Figure 3.11 outlines the way of obtaining the spatial resolution $r_s(d)$ for any given distance d. As a first step, a measurement series $s_n(d)$ with standard variance $\sigma_n(d)$ is linearly interpolated. For each distance d, $\sigma_n(d)$ is applied to analyze distances d_1 and d_2, providing measurement values which differ from $s_n(d)$ within $\pm\sigma_n(d)$. The result is a measure of the spatial resolution $r_s(d)$, which is related to the absolute difference between the two distances. In other words, the spatial resolution is an indicator for the precision $\pm r_s(d)$ a capacitive sensor is able to detect a defined object at a certain distance d. Therefore, a small spatial resolution is desirable in application prototypes.

3.2.4. Performance evaluation and influence of electrode materials

Using r_s as the main characteristic value of a capacitive sensor system, it is reasonable to investigate if different electrode surface materials result in diverging spatial resolutions. Additionally we investigate the dependency between the size of a sensing electrode surface and the resulting spatial resolution, which provides essential information that is required when prototyping capacitive sensor applications. The sensing electrodes test set contains rectangular shaped copper electrodes of various sizes. Furthermore, we compare copper electrodes to electrodes made of ITO and PEDOT:PSS. Using the introduced test setup and measurement method, the normalized sensor characteristics and the spatial resolution were determined.

As the spatial resolution is calculated on interpolated curves of the sensor characteristics and their

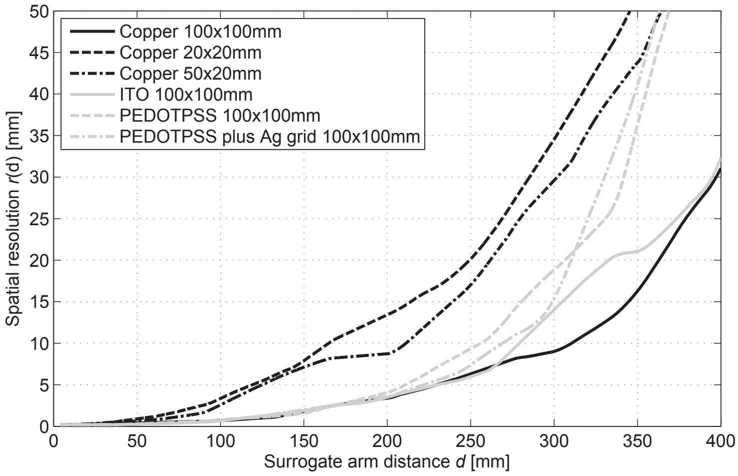

Figure 3.12.: Spatial sensor resolution of a loading mode sensor in relation to the surrogate arm distance.

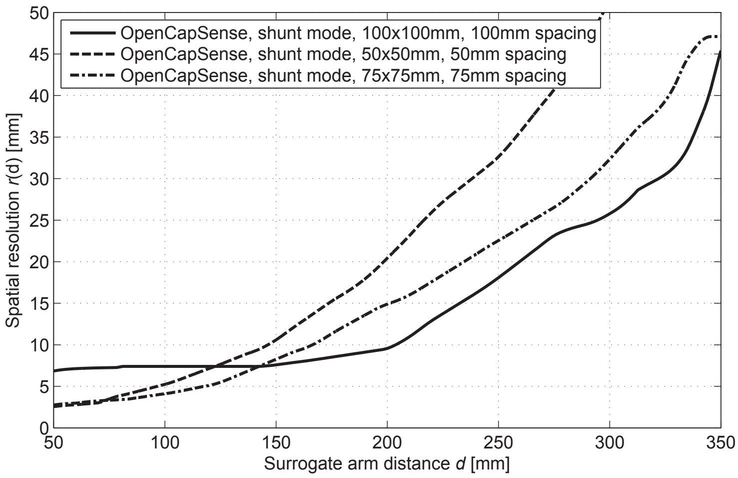

Figure 3.13.: Spatial sensor resolution of a shunt mode sensor in relation to the surrogate arm distance.

corresponding standard deviation, its raw curve is not uniform. Therefore the r_s curves were smoothed, which allows for comparable inferences on the resolution.

Regarding the loading mode sensor, we evaluated the influence of different electrode materials and sizes. The final results for a measurement window of 10 ms are shown in Figure 3.12. As expected, larger electrode sizes provide a better resolution than smaller ones. At distances of 20 cm, the resolution of a 100 x 100mm copper electrode is almost four times higher than the resolution of an electrode with a size of 20 x 20 mm. Regarding different materials, transparent ITO electrodes perform almost as well as copper electrodes. Although ITO has a higher surface resistance than copper, the influence of that property turns out to be of low significance due to the very small displacement current flowing from the electrode to grounded parts in the environment. PEDOT:PSS showed similarly good resolution for near distances. At distances above 30 cm the resolution decreases strongly, which can be attributed to the non-uniform application of the polymeric conductor on the PET caused by inkjet printing. We can conclude that larger electrodes perform significantly better in loading mode when detecting objects at great distances. The transparent properties, easy inkjet printing and the good performance of PEDOT:PSS

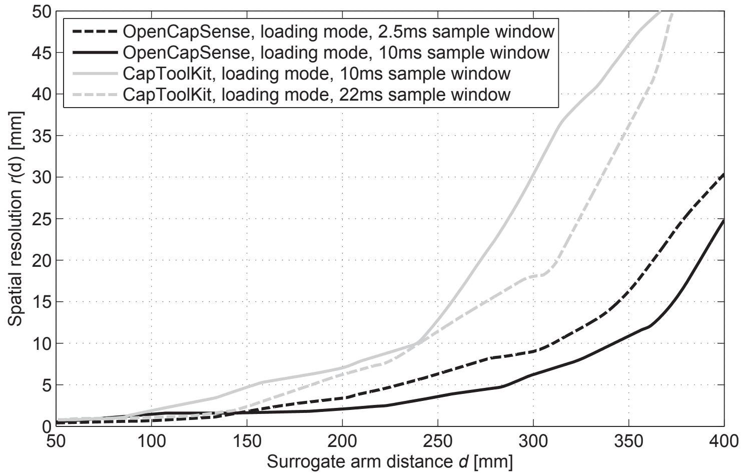

Figure 3.14.: Comparison between spatial resolutions of OpenCapSense and CapToolKit [WKBS07a] using a 100 x 100 mm copper electrode.

qualifies the conductor for rapidly building prototypical user interfaces. In later development stages, ITO can provide a very good performance comparable to copper.

The evaluation results for a shunt mode sensor with a measurement window of 20 ms are illustrated in Figure 3.13. The results show that smaller electrode sizes provide a better performance at low distances, while bigger electrode sizes are more suitable for detecting objects at high distances. It is notable that the measurement object turns into transmit mode at object distances below 5 cm, letting the sensor values increase again. This distance depends on the electrode size: when using small electrodes with a size of 50 x 50 mm, transmit mode comes into effect at approximately 2 cm, while the distance increases to about 5 cm on larger electrodes sized 100 x 100 mm. Although a single shunt mode sensor shows a significantly lower spatial and temporal resolution than a loading mode sensor, it must be recognized that shunt mode offers advantages due to parallel transmitter operation and the exploitation of different receiver-transmitter combinations.

Based on our experiment setup, we evaluated the loading mode sensor for both, CapToolKit and OpenCapSense, using the same methodology. Figure 3.14 shows a comparison of different sampling windows. As expected, the resolution of both, OpenCapsense and CapToolkit, increases for longer sampling windows. Comparing both 10 ms measurement series, OpenCapsense shows a distinct improvement to CapToolkit. For distances of 200 mm, OpenCapSense improves the resolution by the factor three compared to CapToolKit. Regarding a reduced sample window of 2.5 ms the resolution is robust and higher than CapToolKit. As CapToolKit supports minimum measurement windows of 10 ms, a smaller window could not be selected for comparison with OpenCapSense. It can be concluded that OpenCapSense's temporal and spatial resolution is a significant improvement to CapToolKit and allows investigating new pervasive application scenarios with very fast update rates.

3.3. Rapid Prototyping Examples

In order to demonstrate the capabilities of the OpenCapSense toolkit and to observe its properties in real application scenarios, we built a variety of prototypical systems that we will present in the following.

3.3.1. Gesture Recognition

Gesture recognition is one of the mostly applied application scenarios that has been evaluated for capacitive proximity sensors [WKBS07a, ZSP*95, BH09]. The unobtrusive applicability allows detecting gestures over a distance through non-conductive materials. The position of one or more hands can be determined using different analytical or probabilistic models.

Figure 3.15.: CapTap controlling a multimedia application with radial menus. It is based on an array of OpenCapSense boards and an accelerometer to detect knocking and tapping [BZWK*14].

A prevalent issue, present in all publications above, is the speed in which the hand locations are determined - a factor that greatly influences the perceived quality of interaction. Delays that surpass real-time are considered unacceptable [PSH97]. Regarding OpenCapSense's sophisticated processing capabilities, gesture recognition applications can be implemented on the demonstration board itself. Combined with a high sampling rate this enables real-time hand gesture recognition while preserving a good precision. A recent prototype we are developing is the CapTap - a combined hand tracking and knock detection system that uses multiple OpenCapSense boards configured as a sensing array and a single accelerometer interfaced through I^2C. The system, shown in Figure 9, allows controlling typical multimedia applications with selection indicated by hand position and actions triggered by different knocking events that are registered by the accelerometer.

3.3.2. Building Automation

By now, capacitive sensing has not yet pervaded the field of home or building automation. One reason can be seen in the great variance of possible physical setups and thus the low amount of generalizability. Especially when deployed underneath a floor, capacitive sensors in home-automation must support a wide amount of load capacitances, which can range from 100 pF to 1 nF. After OpenCapSense was released, I was asked for consultation in an interactive art installation at San Francisco Airport. The installation, shown in Figure 3.16, reacts on persons touching the circles by adapting their light and motion.

One key requirement was to make the sensor compatible to a programmable logic controller (PLC). Therefore, I extended OpenCapSense's loading mode sensors by a microcontroller and two optocouplers, that can switch the PLC's 24 V signal. Another optocoupler is used for resetting the capacitive sensor. The two outgoing signals indicate either touch and proximity, by analyzing the measured capacitance. Therefore, copper electrodes are attached besides the moving illuminated plates. A major challenge in the installation has been the large static capacitance.

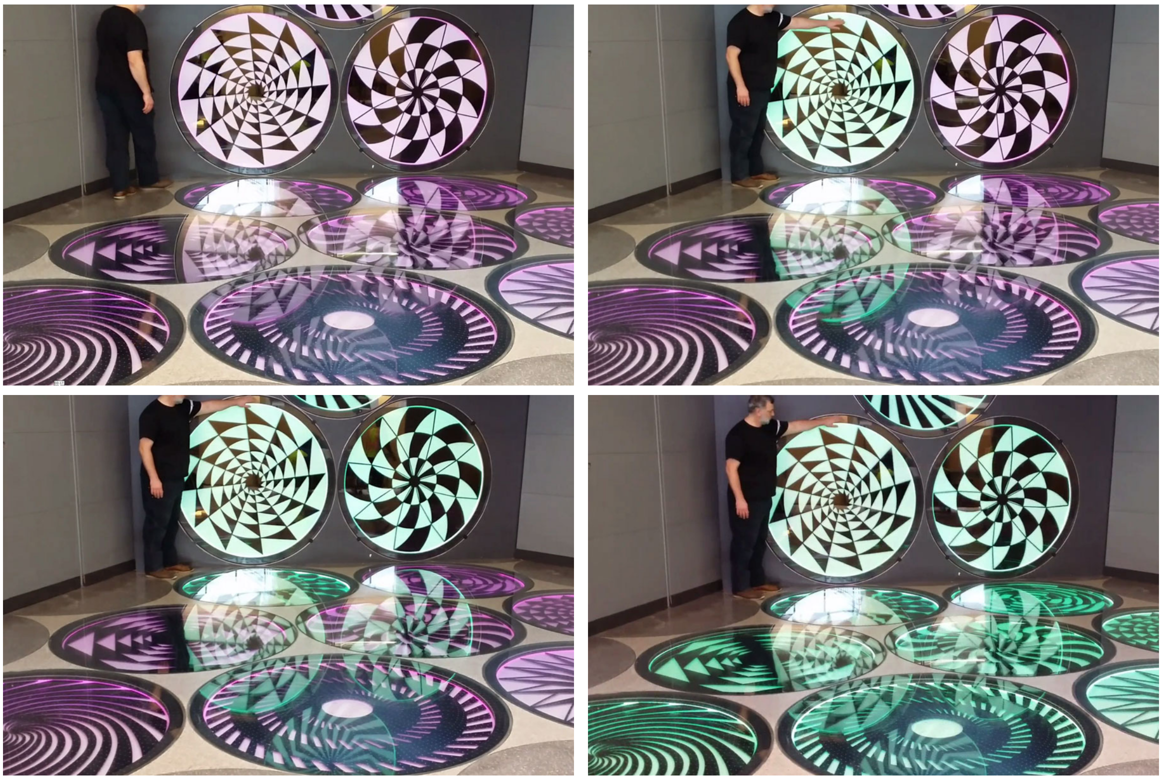

Figure 3.16.: Spirogyrate is an interactive art installation at San Francisco Airport by Eric Staller [Sta14], realized by RocketScience, San Francisco. Here, a modified version of OpenCapSense is used to provide interactive capabilities.

Figure 3.17 shows the sensor values obtained from the OpenCapSense sensor for different experimental setups. Due to the large static capacitance, the sensing noise was mostly in the region between +/- 1. The measurements took 100 ms for still being able to measure a reasonable change in signal when body parts are present. Interestingly, without glass placed above the electrodes, the delta between no presence and presence is very low. With glass above the sensing electrode, the delta largely increases although the base capacitance becomes bigger. In order to avoid static charges on the glass, we chose to install antistatic plexiglass. Due to the material's conductivity, which is in the region of a few $M\Omega$s, it acts as an electrode resulting in a better application performance. Finally, RocketScience decided to use 3" shields and 2" sensing electrodes.

3.3.3. Fall Detection

Many pervasive environments require affordable and accurate methods for tracking humans. Capacitive proximity sensors are especially suited for this use case as they can be integrated unobtrusively into the environment. In real use-cases, this monitoring requires to strongly consider privacy aspects. Researchers have therefore realized floor tiles and carpets employing capacitive proximity sensing that can be deployed in assistive environments and home automation scenarios [SL08, VMV09]. Smart floors can support the elderly and disabled persons in their activities of daily living, for example by monitoring emergency situations such as falls.

Current capacitive proximity sensing systems for localization and fall detection face the problem of

Setup	No Foot	Foot Present	Delta	Relative Delta
3" width shield, 0.5" width sensor				
no glass, no ground	2165	2154	-11	-0.5%
no glass, w/ ground	2147	2128	-19	-0.9%
w/ glass, no ground	1838	1764	-74	-4.0%
w/ glass, w/ ground	1800	1725	-75	-4.2%
3" width shield, 2" width sensor				
no glass, no ground	1370	1342	-11	-0.5%
no glass, w/ ground	1373	1348	-19	-0.9%
w/ glass, no ground	1031	989	-42	-4.1%
w/ glass, w/ ground	1013	964	-49	-4.8%

Figure 3.17.: Evaluation of different electrode setups conducted by RocketScience, San Francisco. The measured sensor values refer to the number of charging and discharging cycles within 100 ms.

having a very low update rate of 10Hz or less [VMV09]. Thus, a fall is usually detected by evaluating the time a person lies on the floor. However, a fall detection could benefit greatly from a very high update rate that enables a system to reconstruct the course of the fall precisely. We placed eight sensors and their corresponding electrodes (30 x 20 cm) under an ordinary carpet. The smart carpet can detect a fall with a very high update rate of approximately 104 Hz for each sensor using a round-robin sensor scheduling. Therefore, we have reduced the measurement interval to 1.2 ms. This very small measurement interval is associated with increased noise but can reliably detect a fall situation as shown in Figure 3.18.

3.4. Summary

With OpenCapSense, I am able to provide an easy-to-use prototyping environment for UbiComp and HCI applications. The work is a scientific contribution to the first research challenge on new capacitive sensing approaches. OpenCapSense covers four of five dimensions in the domain of proxemic interactions. It provides the physical sensing opportunities to measure distance, orientation, movement and location. The remaining dimension of sensing the identity of objects will be tackled in the next chapter.

OpenCapSense is the first prototyping toolkit that supports all proximity sensing modes, namely shunt-mode, transmit-mode and loading-mode. The toolkit provides eight generic channels for connecting different types of sensors. It offers sophisticated processing capabilities that support real-time FFTs at sampling rates of up to 100 kHz.

The toolkit was evaluated for spatial resolution, the influence of different electrode materials on this metric, and comparatively analyzed with the CapToolKit [WKBS07a] - a platform that is used in other research projects. Developers are able to easily connect the board to other devices by providing interfaces for USB, CAN, and I^2C. After publishing the main paper about OpenCapSense [GPBB*13], we also added an interface for Bluetooth communications. This is especially valuable when integrating the toolkit in movable furniture, such as office chairs.

I validated the applicability of OpenCapSense by presenting a number of different case studies, ranging from light wearable devices to fall detection systems in fixed installations. Still, these applications represent only a small subset of potential scenarios that can be realized. After almost two years of existence, OpenCapSense has been applied in various scientific works and thesis [Ber12, BZWK*14, BNS*14, DBM14, GPMB11, GPBB12, GPBK*13, Rus13, RGPK14, ZMB*14]. Moreover, in the European project *Flexible Smart Spaces*, it is employed in posture-sensing office chairs that have been distributed to secretaries in Finland. The chair helps secretaries to maintain personal fitness by giving them information on how to improve the seating posture or performing certain exercises as an office workout.

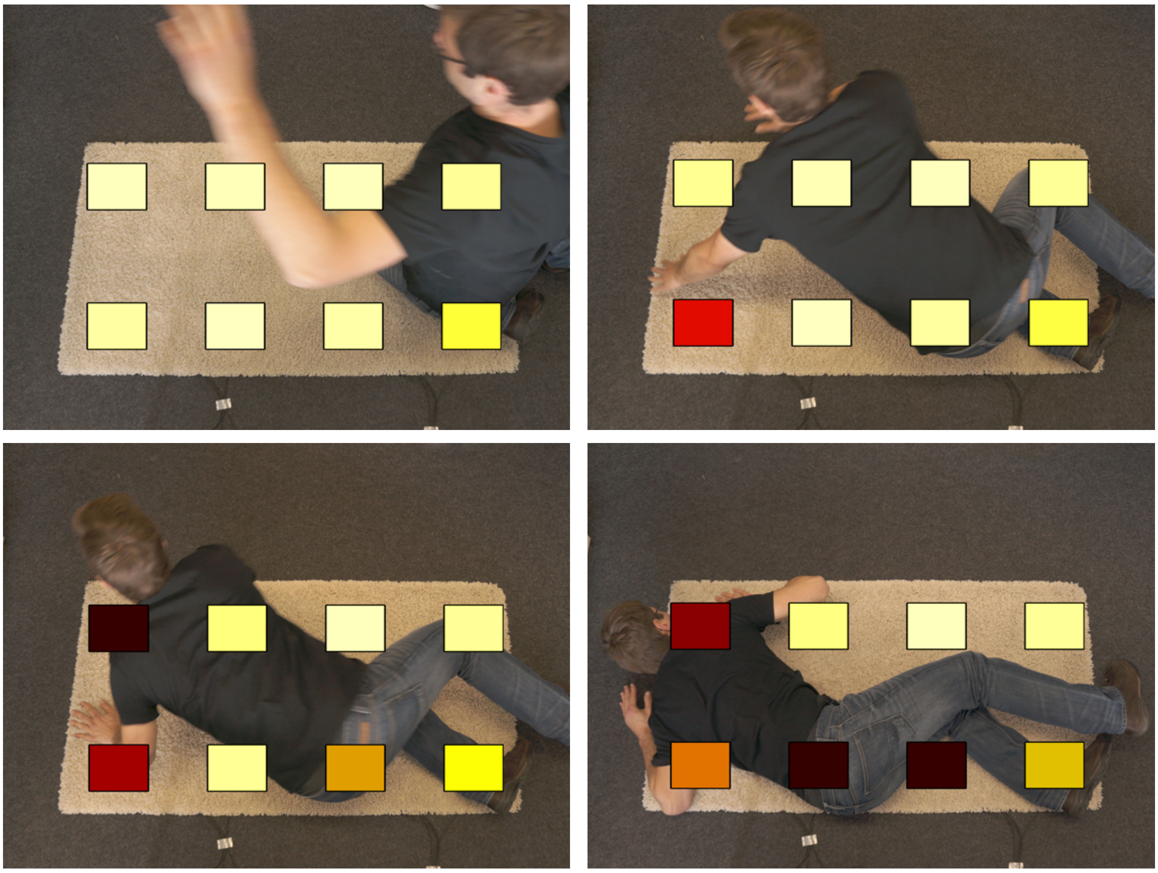

Figure 3.18.: Fall detection can greatly benefit from very high sensor update rate (1.2 ms measurement window per sensor): red sensor values denote near measured distances and white sensor values represent no body parts in the proximity of the sensor. The simulated fall situation lasted approximately 800 ms.

In the European research project *POSEIDON*, OpenCapSense is deployed in a prototypical interaction device for people with Down's syndrome [Bra14]. Here, the goal is to give support in structuring the daily routine and help to learn new concepts and reinforce the learned contents.

4. Capacitive Near-Field Communication

In the previous chapter, I have shown how to cover a subset of proxemic interaction dimensions introduced in [GMB*11]. Based on a novel capacitive sensing toolkit, I am able to provide physical measures for the dimensions of distance, orientation, movement, and location. However, there is still one dimension left which I could not approach with pure capacitive sensing: Measuring identification. Identification is not limited to recognize the individual object, but also to identify its inner state. This state is not perceivable by external sensors and can comprise much more fine-grained information. Although some researchers have integrated soft biometrics in capacitive touch sensing [HSP12], identification of objects and humans is still an open research question in capacitive sensing.

Besides measuring orientation and identification, it is necessary to revise the dimension of measuring distance. I think that we need to move away from a quantitative view of distance to a more qualitative notion. The reason is that proxemic interactions depend highly on the context of a human being. For example, two objects can be regarded as close to each other when a person touches them. Still, they might be located some distance away from each other, but the human action tightens the connection between both. This example underlines that the notion of distance between objects must be weakened. A human-centric view on this dimension is necessary to achieve more intelligent interaction concepts.

In order to measure the missing dimension of identification (shown in Figure 4.1), I introduce a novel communication infrastructure based on capacitive coupling. The idea is based on mutual collaboration: intelligent tags can be attached to objects and humans. The tags are able to measure information about their environment, for example by employing inertial sensors. Based on a broadcasting principle, they communicate their manipulations to other objects in range employing capacitive coupling. The refined dimension of distance is approached by allowing the tags to communicate through the human body as well as through air.

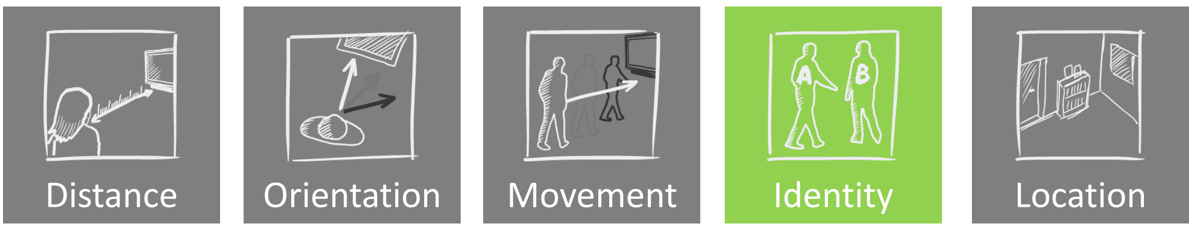

Figure 4.1.: Proxemic interaction dimensions achieved with Capacitive Near-Field Communication [GMB*11].
[GMB*11] © 2011 Association for Computing Machinery, Inc. Reprinted by permission

Using capacitive coupling for communication was first applied by Zimmerman, who investigated personal area networks with nodes communicating through the human body [Zim96]. Moreover, by means of capacitive coupling, information has been exchanged between smart artifacts and touchscreens. In this chapter, I generalize over previous approaches and present a common notion for capacitive near-field communication. Based on this generalized notion, I demonstrate the communication method's applicability by three application examples.

This chapter is based on the paper [GPHW*14], which includes evaluation results from Sebastian Her-

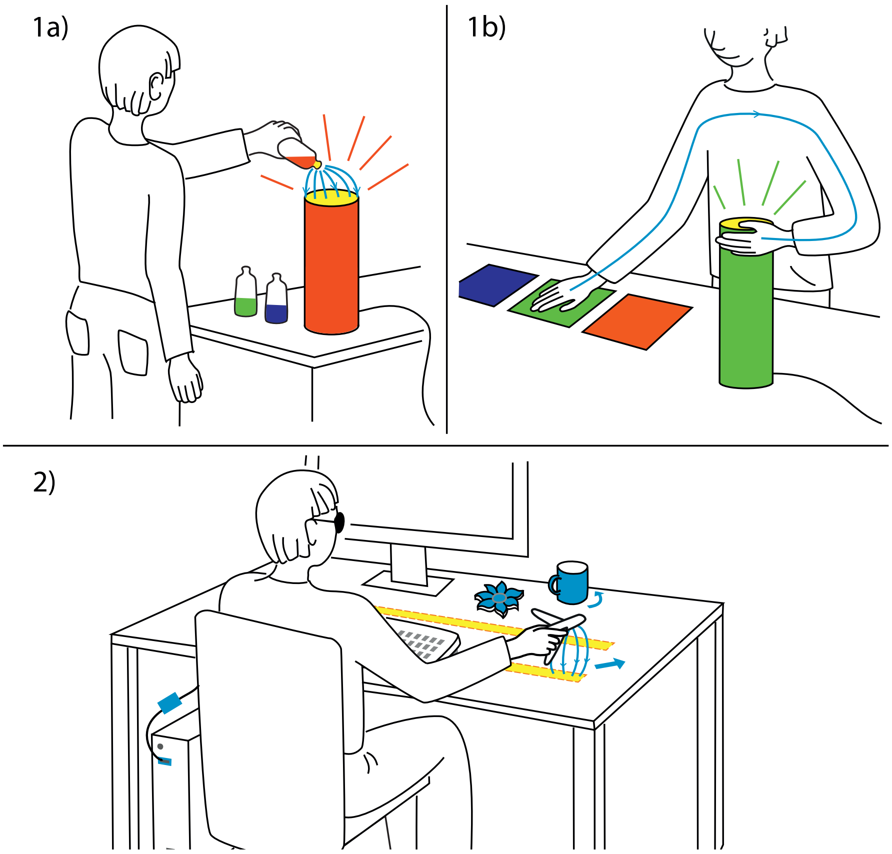

Figure 4.2.: CapNFC use-cases: (1a/b) Interacting naturally with everyday objects by transmitting object-related information through air or the human body, (2) using everyday objects to support blind users in interacting with a computer.

ber's Bachelor's thesis [Her13]. The use of '*we*' refers to the paper's authors Tobias Grosse-Puppendahl, Sebastian Herber, Raphael Wimmer, Frank Englert, Sebastian Beck, Julian von Wilmsdorff, Reiner Wichert, and Arjan Kuijper. Our work puts a new focus on ubiquitous interaction with a more generalized discussion of the communication technique and its properties. Therefore, we offer several scientific contributions in this chapter:

1. We introduce and evaluate a novel conceptual basis for Capacitive Near-Field Communications in ubiquitous environments. We identify a set of operating modes and the corresponding interaction techniques.
2. Three case studies support the applicability of CapNFC in Ubiquitous Computing: (i) natural interaction paradigms with everyday devices, (ii) smart objects assisting blind users in interacting with their computer, and (iii) a smart bed with wearable sensors for sleeping behavior analysis. The first two case studies are partly depicted in Figure 4.2.
3. We describe, compare and evaluate the concepts of CapNFC in sufficient detail to allow others to re-implement it. Open-source schematics and source code enables the community to quickly implement CapNFC applications.

4.1. Motivation

Technologies like RFID or NFC represent very suitable technologies in communication when the information is exchanged in a short-range spatial context. Moving from the inductive to the capacitive domain, a human's conductive properties allow for indirectly incorporating the perception of body parts, touches and proximity into the communication technique. Especially in ubiquitous computing, this property leverages its full potential as a means for short-range communications using the air or the human body as a communication channel [Zim96].

In the following, we present and discuss Capacitive Near-Field Communication - CapNFC - as an ultra low-power, touch- or proximity-mediated alternative to existing wireless communication techniques. Existing smart objects can be easily adapted to use CapNFC as only a single microcontroller output pin is required for unidirectional communication. In terms of form factors, CapNFC is extremely flexible, supporting the integration into beds as well as small objects, such as body-worn sensors. CapNFC employs capacitive coupling between nearby conductive electrodes to transmit information over distances of up to 15 cm through air. Furthermore, the technology allows for communicating messages through the human body or measuring the proximity to human body parts. It utilizes frequencies of 1 kHz to 1 Mhz, with wavelengths that are a multitude larger than the transmitting and receiving electrodes. Therefore, CapNFC acts in the quasi-electrostatic regime, eliminating the necessity for high-frequency design and complex hardware components. This results in a very small displacement current flowing from a transmitter to a receiver electrode, enabling ultra-low power operation.

4.2. Conceptual Basis for Capacitive Near-Field Communication

CapNFC relies on *capacitive coupling* between a transmitter and a receiver electrode [ZSP*95]. Applying an alternating voltage to one electrode and connecting the second electrode to ground causes a changing electric field between both electrodes. With CapNFC, the induced *displacement current* is measured on the receiver side. The energy-saving potential in ubiquitous communications can be regarded by physical considerations. CapNFC acts in the quasi-electrostatic regime and thus without any wave-propagation at its electrodes. This property solely enables a very small displacement current to flow from the transmitting electrode to the environment. In order to communicate by means of capacitive coupling, voltages oscillating at low frequencies ($\leq$ 1 MHz) are applied to the transmitter electrode. The receiver amplifies the induced displacement current and decodes the message. Various encoding schemes may be used on top of this physical communication channel. Communication range and robustness depend strongly on the size of the transmitter electrodes and their voltages levels. Higher voltages increase the displacement current and lead to a better performance. Bigger electrodes allow for better capacitive coupling [GPBB*13] but also pick up more background noise. In our examples, electrode sizes vary from 1 cm^2 to 100 cm^2.

In this work, we apply the concept of CapNFC to ubiquitous interaction with everyday objects. Each object which is touched or manipulated by a user establishes an information flow and communicates its state to all objects within range. Devices with extended interaction capabilities, such as smartphones and computers, can interpret the information and generate a response to object manipulations in their environment. In the following, we discuss general operating modes, describe our reference implementation comprising hardware and communication protocol, present considerations for electrode placement, and document performance and energy consumption of our prototypes.

4.2.1. Operating modes

All communication techniques based on capacitive coupling do not only require a transmitter and receiver electrode, but also a connection to a common ground. Only when two objects share a common ground, a displacement current is able to flow from one electrode to the other. A common ground may be established by sharing a ground between both circuits or by connecting each circuit to the environment's ground. However, as only very small displacement currents flow back and forth, CapNFC also works if a transmitter and a receiver are grounded via weak capacitive coupling to the environment's ground. This connection can be generated by a user touching the object, or a table on which the object is placed.

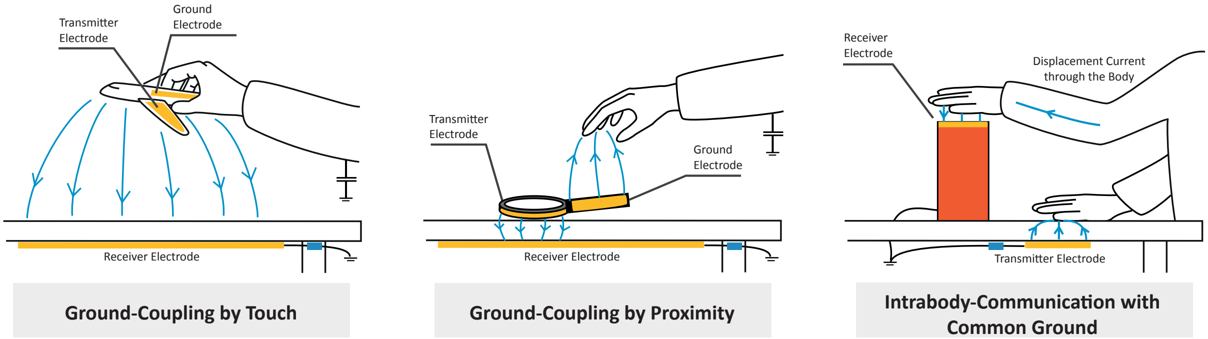

Figure 4.3.: CapNFC offers three operating modes that depend on the ground-coupling of objects: Ground-coupling by touch for indirect touch recognition, ground-coupling by proximity for sensing a person approaching an object, and sharing a common ground for intrabody-communication.

Many ubiquitous artifacts and devices are portable and must thus be battery-powered. Therefore, they lack a shared ground or direct connection to the ground, which is a pre-requisite for communication. Nevertheless, humans usually have good capacitive coupling to the environment's ground [Zim96]. Thus, establishing a common ground between two objects can be achieved by embedding ground electrodes into the casing of an object and connecting the electrode to the battery's negative terminal. When a human touches the ground electrode or is in close proximity to it, capacitive coupling between the ground electrode and the environment's ground is sufficiently strong to allow a displacement current to flow. A very strong coupling to the environment's ground, for example by directly connecting the objects to a grounded power supply, enables objects to communicate messages through the human body [Zim96]. Based on these properties, we identified three different operating modes for CapNFC that allow for reliable data communication in ubiquitous environments, as depicted in Figure 4.3. CapNFC embraces and utilizes the need for a common ground connection between sender and receiver by making it a central part of the supported interaction methods.

Ground-Coupling by Touch: In this mode, the receiver object is connected to the environment's ground, and the transmitter is a battery-powered object. Besides the transmit electrode located in the smart object, a small additional electrode is connected to the battery's negative power supply. Whenever the user is in very close proximity to this electrode or touches it, the weak capacitive coupling to the grounded human body enables an information flow from the transmitter to the receiver. Therefore, the ground electrode indirectly acts as capacitive touch sensor.

Ground-Coupling by Proximity: In this operating mode, a receiver is connected to the environment's ground, and a battery-powered smart object is located in close proximity to the receiver. In addition to the transmitter electrode, the battery-powered object has a built-in ground electrode. When the user approaches this electrode with a body part the capacitive coupling to the environment's ground increases.

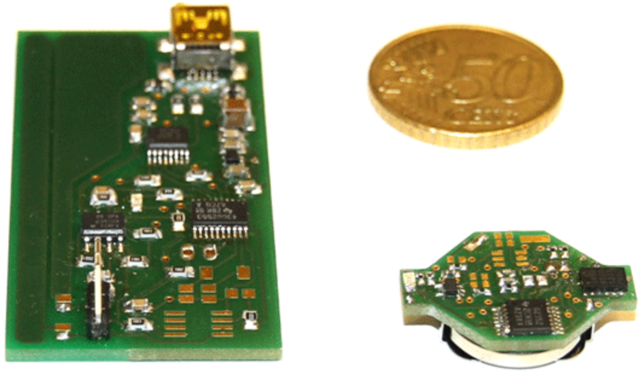

Figure 4.4.: Two different CapNFC components: A tag (right) with an embedded accelerometer and a transceiver (left). The hardware is only partly equipped to allow for experimenting with active sending filters.

By measuring the signal strength in the receiver, it is possible to recognize an approaching human hand in distances up to 15 cm.

Intrabody-Communication with a Common Ground: When both objects are directly connected via a common ground, an information flow can be established through the user's body [Zim96]. This enables the user to touch two objects and thus establish communication between both of them. In CapNFC, this mode is primarily applicable for stationary installed objects, as they may be directly connected to a grounded power supply.

4.2.2. Reference Implementation

The CapNFC concept can be adapted for a multitude of applications. However, we believe that ubiquitous user interfaces benefit the most from CapNFC. Before designing the CapNFC reference hardware and software, we defined constraints that must be fulfilled in order to allow for a high number of low-cost smart objects: (1) transmitting information should not require any additional active hardware components (crystals, amplifiers, ...) except for a single microcontroller; (2) receiving information must consume a low amount of processing power and energy; (3) the method has to be suitable for low-power microcontrollers that run on frequencies as low as 1 MHz. For Ubiquitous Computing, we prioritize these goals over a high data rate.

Based on these design-centered decisions, we implemented two hardware components: a transceiver and a low-power transmitter tag, as shown in Figure 4.4. The transceiver comprises a microcontroller (MSP430G2553), a serial-to-USB converter (FT230XS), an operational amplifier IC (OPA2344), and a voltage regulator (TPS79733). It can be connected to a PC or smart phone via USB. The tag contains only a low-power microcontroller (MSP430G2352) and an accelerometer (ADXL345). It is powered by a 3 Volt coin cell monitored by a voltage regulator (TPS79730). The tag can be attached to an object and is only able to transmit pre-defined and context-dependent messages via capacitive coupling, using all three modes identified in the previous section. The accelerometer allows saving power and supporting gestural interaction with a tagged object. It activates CapNFC communication whenever the tag is moved or a pre-defined gesture is detected.

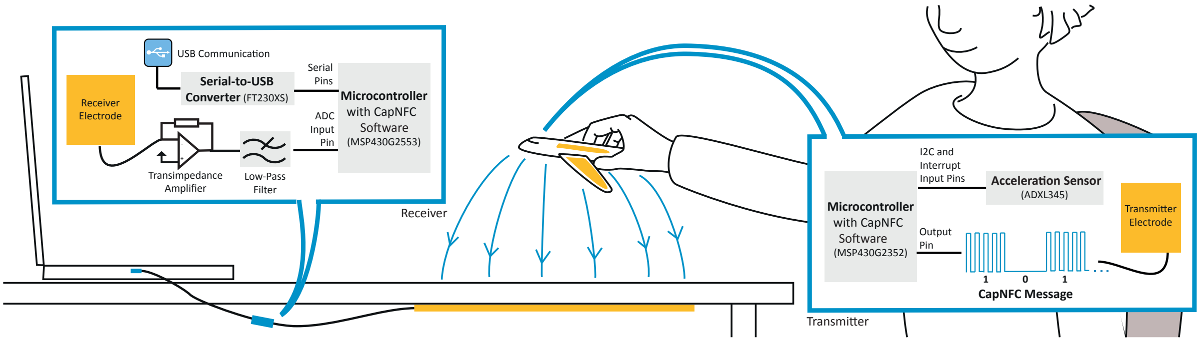

Figure 4.5.: An exemplary CapNFC setup: A toy airplane communicates with a computer by transmitting its movement related data using a quasi-electrostatic field. The airplane is equipped with a low-power CapNFC tag, consuming 280 μA in total while transmitting a message. The receiver amplifies the induced displacement current, decodes the messages and hands them to the computer via USB.

4.2.3. Communication Method and Protocol

Figure 4.5 shows an exemplary CapNFC communication setup. The communication principle itself is based on broadcasting messages; the range around an object can be represented by a spatial context on which messages of predefined types are published. Components may also publish higher level information composed of previously received as well as data measured by themselves.

Due to the design constraints and the limited communication space, we decided not to include carrier-sense multiple access techniques (CSMA) to avoid collisions when multiple transmitters access the communication medium. Instead, the existence of data transmission collisions is accepted deliberately in our protocol. This fact enables smart objects that solely employ a low-cost microcontroller without additional hardware. These objects are only capable of transmitting messages, without being able to receive any information from other objects. Due to this simple concept, existing smart objects can be easily extended to support CapNFC communication as only a single additional output pin at a microcontroller is needed.

Start 10 Bit	Data Length 8 Bit	Type 8 Bit	Sender 8 Bit	Data 0 – 20 Bytes		
Sync	6 Byte	Inertial Sensor	Sender Address of Magnifier Glass	Acceleration x-axis (2 Byte)	Acceleration y-axis (2 Byte)	Acceleration z-axis (2 Byte)
0x03 0xAA	6	0x01	0x12	10	-15	140
Sync	3 Byte	Color Picker	Sender Address of Mobile Phone	Red (1 Byte)	Green (1 Byte)	Blue (1 Byte)
0x03 0xAA	3	0x20	0x20	255	50	50

Figure 4.6.: Our CapNFC communication protocol. The message format can be seen on top, two exemplary messages below. All messages are secured with parity bytes for bit error correction.

Information is transmitted via a 3.0-3.3 V rectangular-shaped signal, oscillating at a frequency of 10 kHz. The signal can be easily created using a microcontroller's pulse-width-modulation module, which is a common functionality even in low-cost hardware. In order to represent binary information, the signal is switched on and off (On-Off-Keying), allowing for a data rate of 2 KBit/s. As illustrated in Figure 4.6, a message begins with two high bits which represent a start condition, followed by a synchronization byte of succeeding 1-0 transitions (0xAA), the length of the message (1 byte), the message type (1 byte), a sender address and a succeeding payload (less than 30 bytes including parity bytes). The reason for the payload length limitation lies in potential clock-drift among the different hardware components. The receiver of a message synchronizes with each 1-0 and 0-1 transition. However, when the number of transitions is low, the clock-drift would be too high for payloads longer than 30 bytes.

In order to receive data, which is optional in CapNFC, a receiver employs a transimpedance amplifier followed by a low-pass filter [SGB99]. It requires three simple hardware components - a low-cost operational amplifier, two resistors and a capacitor. The receiver's microcontroller samples the signal at 40 kHz, and reconstructs the transmitted information with the Goertzel algorithm [Goe58]. This computationally inexpensive algorithm is used to detect the presence of the transmitted signal oscillating at 10 kHz. In the following, parity bytes are used to restore potentially corrupted bits and an optional CRC checksum ensures the integrity of a message.

Figure 4.7.: Different electrode placements: (1) The transmitter electrode is placed directly under the bottle cap, (2) the magnifying glass has a big ground electrode in the region where the user touches the object, the transmitter electrode is placed around the glass, (3) the lighter uses its conductive part as a transmitter electrode, and (4) the mobile phone's transmitter electrode is placed at its back, the lamp's receiver electrode is under the black surface.

4.2.4. Electrode Placements

As stated in the introduction of CapNFC operating modes, communication between two objects requires a weak capacitive coupling to a common ground. Besides connecting the object directly to a common ground, a weak ground coupling can be set up by a human touching or approaching a ground electrode. In order to establish a suitable coupling to the environment, this electrode is attached to the object's negative supply voltage and placed near the object's surface or directly connected to conductive parts of the object. In his work about Personal Area Networks, Zimmerman determined a capacitance of 110 pF [Zim96] for a current flowing from an isolated human body to the environment's ground.

If the object is only supposed to communicate while being touched or when a hand is in proximity, it is necessary to place a small electrode somewhere near the regions the hand approaches. The bigger the ground electrode gets, the better a displacement current is able to flow from the transmitter electrode to a different object's receiver electrode. Transmitting information without physical touch can be realized by placing an electrode close to other objects having an indirect connection to the environment's common ground, for example a table's wooden surface. The ground electrodes may also be isolated or located within the object as it is only necessary that alternating current is able to flow. In the easiest scenario, the transmitting or receiving electrodes of a CapNFC-enabled object are the object's conductive body itself. However, electrodes can be constituted of different materials, for example conductive threads, aluminum foil, transparent ITO layers or inkjet-printed silver [KHC*13]. The low surface conductivity of transparent or printed materials themselves decreases the performance of capacitive coupling applications only slightly [GPBB*13]. This fact can also be transferred to capacitive communications and enables a designer to choose from a wide variety of materials with different properties. Figure 4.7 shows possible electrode locations and placements of some exemplary CapNFC components.

4.2.5. Performance Evaluation

We evaluated the performance of the presented communication method by assessing our reference implementation with different tagged objects having uni- and bidirectional communication abilities. The experiments were carried out in a standard living environment, with the objects being placed on a wooden desk. As a performance indicator, the signal-to-noise ratio ($\text{SNR} = 10\log_{10}(P_{\text{signal}}/P_{\text{noise}})$) was applied. In the following, we present evaluation results for all three CapNFC operating modes.

Figure 4.8 depicts the SNR for different transmitter/receiver distances employing *ground-coupling by touch*. Therefore, we combined two copper electrodes of different sizes and cross-evaluated the SNR. Interestingly, the evaluation revealed that one should rather use larger transmitter electrodes, than enlarging the size of the receiver electrode. Enlarging the receiver electrode's size will result in a greater amount of noise, whereas the strength of the received signal does not increase proportionally. This is a negative property for ubiquitous interaction, as it is very difficult to increase the transmitter electrode's size in small battery-powered objects. Enlarging the receiver electrode's size, which may be located under the desk, is usually easier as the space is not as constrained. Depending on the electrodes' sizes, messages can be transmitted in hand distances up to 15 cm.

Figure 4.9 depicts the evaluation results for *ground-coupling by proximity*. In this experiment, we use a battery-powered tag with differently sized ground and transmitter electrodes, located in close distance to a receiver with a corresponding large receiver electrode. A human hand was used to approach the ground electrode at different distances with multiple repetitions of the experiment. The experiment shows that the detection distance increases with larger ground electrodes. The evaluation of larger ground electrodes revealed that the weak coupling to the environment's ground, for example the table, is sufficient for transmitting messages. In these cases, the SNR remains constant even after removing the hand. Alas, inferring a certain distance from a given SNR requires strong initial assumptions about the object's location and orientation. Nevertheless, differences in the SNR between succeeding messages can be exploited to detect hand presence and proximity differences.

The evaluation results for *intrabody with common ground* are depicted in Figure 4.10. We evaluated

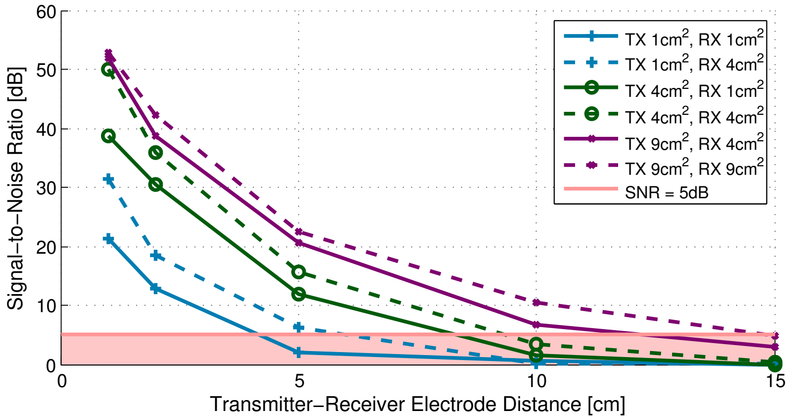

Figure 4.8.: Ground-Coupling by Touch: The Signal-To-Noise Ratio with different receiver/transmitter distances and electrode sizes. The 5 dB SNR threshold results in a bit error rate (BER) of more than 20%, which means that many messages are corrupted.

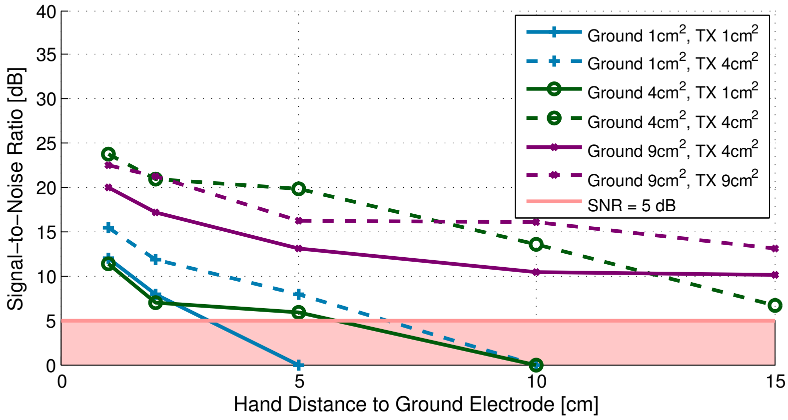

Figure 4.9.: Ground-Coupling by Proximity: The object is placed in little distance to the receiver electrode. The diagram shows the SNR for different ground and transmitter electrode sizes. By approaching the hand to the ground electrode, the SNR increases.

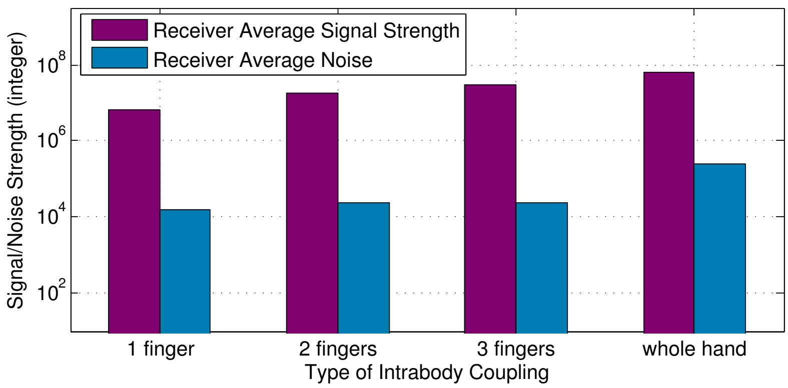

Figure 4.10.: Intrabody with Common Ground: The signal/noise strength is given as an integer value, representing the average result after computing a window with the Goertzel algorithm.

different touch types for coupling a transmitter and a receiver electrode. Therefore, we isolated each electrode under a wooden surface and touched it with one, two, three fingers, and the whole hand. Surprisingly to us, the SNR was not affected very strongly when touching the electrodes with more fingers. However, when analyzing the signal and the noise levels separately from each other, the signal as well as the noise levels increase when the coupling is made through a larger area, e.g. a hand. This interesting property is very useful for gesture or grasp recognition, for example by remotely detecting when a finger was added or removed from an object.

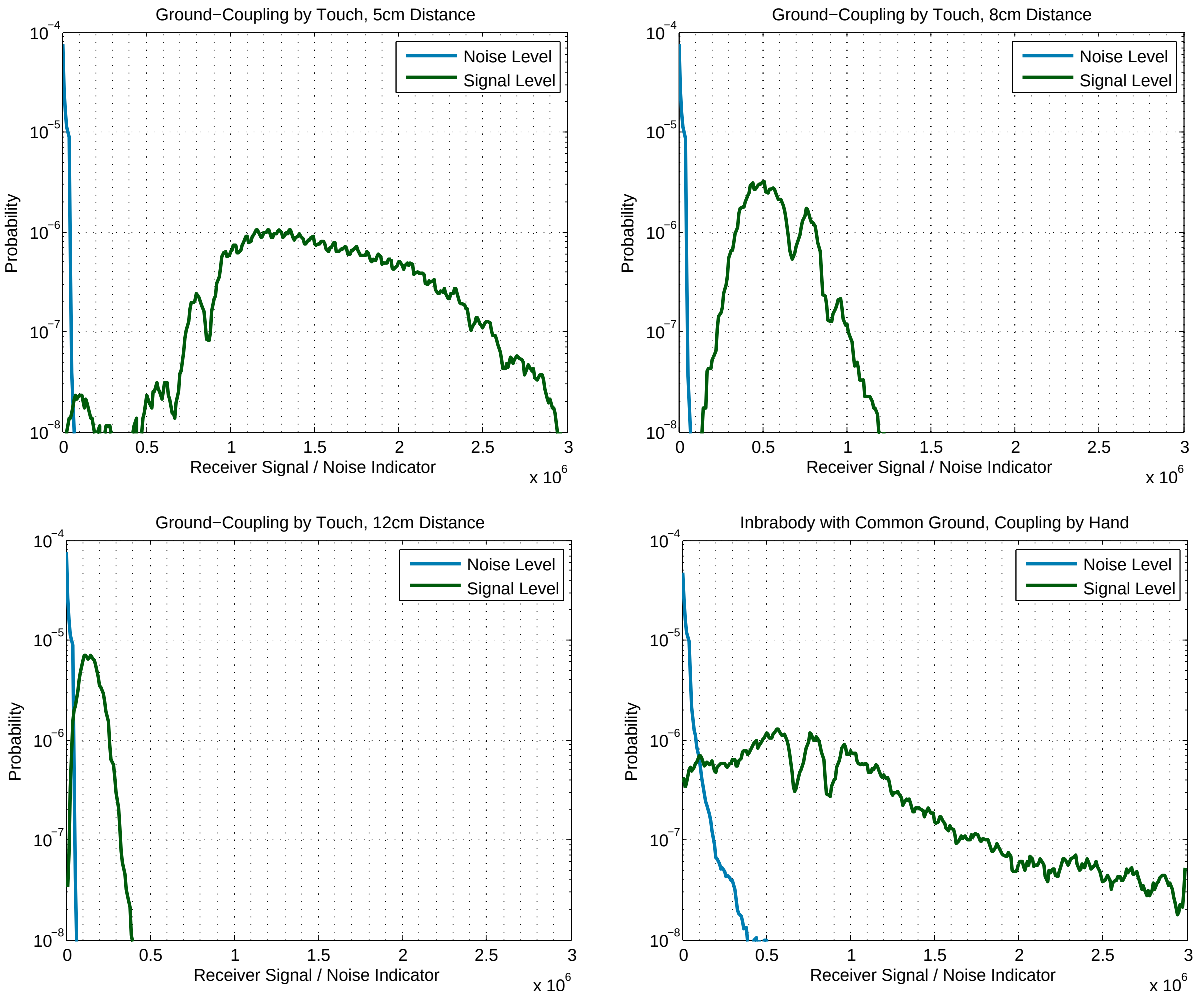

Figure 4.11.: The plots show the probability density for signal and noise based on the received signal/noise strength indicator.

Figure 4.11 depicts the noise and signal strength measured at the receiver. It shows four probability densities to illustrate the impact of different transmitter-receiver distances and intrabody communication. The optimal bit decision boundary can be approximated by the intersections between signal and noise density. When regarding one operating mode only, e.g. ground-coupling by touch, it is very easy to find a suitable bit decision boundary based on estimating the noise. Problems arise as soon as ground-coupling by touch is combined with intrabody communication. In this case, the noise will suddenly increase and the bit decision boundary must be set to a different level. Generalizing over different operating modes therefore introduces the problem of dynamically estimating noise and setting bit decision boundaries.

Figure 4.12 shows the relation between SNR, bit error rate (BER) and packet error rate (PER) in our

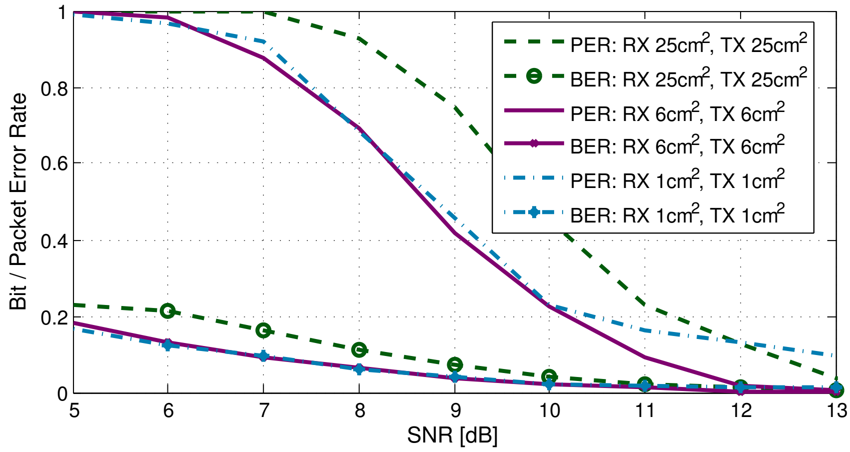

Figure 4.12.: The relation between SNR, the bit and packet error rate in the region from 5 -13 dB. The packets had a data length of 3 byte (+1 byte sync) secured with 3 parity bytes ((8,4)-Hamming).

reference implementation. The PER does not only depend on the BER after error correction, but also on the detection of the message preamble, which is currently not fault tolerant. When using larger receiver electrodes, the already high noise variance increases and leads to a decreased BER and PER. Possible reasons can be seen in noise generated by devices operating at similar frequencies, e.g. switch-mode power supplies.

Due to the currently high noise variance, high SNRs are required to obtain a good packet error rate. Besides lowering error rates, more noise-resilient communication methods could also extend the interaction range to 30 cm. In particular, we plan to apply frequencies greater than 100 KHz and use synchronous undersampling to restore the signal [SGB99]. Increasing the carrier frequency will make the signal less prone to noise produced by other electronic devices. Moreover, we will investigate more sophisticated modulation methods like spread-spectrum modulation to avoid problems with narrowband noise [SGB99].

In the following, we discuss the energy consumption of CapNFC and our reference implementation. The energy required for succeedingly loading and unloading the capacitance between transmitter electrode and grounded parts has only little influence on the implementation's energy consumption. This fact is very important, as it only adds little energy consumption to existing systems while leveraging communication abilities based on capacitive coupling. Moreover, in the future this property allows for creating tags powered by energy-harvesting or capacitive power transfer [Mur11], very similar to RFID. Basically, only a very simple low-power microcontroller is necessary to transmit data. For example, while transmitting 10 inertial samples per second, our low power tag consumes less than 88 μA at 3.0 V on average. Receiving messages requires an additional operational amplifier, which consumes 750 μA in our implementation and a more powerful microcontroller for signal processing. However, the energy consumption of a receiver can be reduced further by choosing a different amplifier. Currently, our amplifier supports a very high slew rate to experiment with significantly higher frequencies than 10 kHz. In terms of energy consumption, it can be concluded that CapNFC supports extremely low-power implementations, especially suitable for the usage in smart, energy-constrained, objects. A comparison of CapNFC with other technologies, such as RFID, will be given in the next section.

4.3. Related Work and Competing Technologies

4.3.1. Capacitive Communications and Sensing

Using capacitive coupling for information exchange is applied in a wide variety of application scenarios. Amongst others, it is used for communication between objects and multitouch enabled devices [VG13, YCL*11], and communication via the human body [Zim96]. Prototyping capacitive sensing applications has made significant advances in the last years due to inkjet-based printing [CJC*06, KHC*13] or low-cost foil cutters [SZH12].

Capacitive Coupling through Air and Direct Contact: Some prototypes of tangible objects allow sending short messages to a device with a capacitive touchscreen by simulating touch events [VG13, YCL*11]. The presented communication technique works with arbitrary capacitive touchscreens and requires only simple hardware for the sender that is integrated into a finger ring. Obviously, the communication is unidirectional from the sender to the multi-touch device and has a low bandwidth limited by the scanning frequency of the touchscreen. In the domain of indoor user localization, Valtonen et al. use conductive tiles on which the human body acts as a transmitter electrode [VMV09]. A receiver placed on the ceiling or within a wall identifies the tiles a user steps on. Cohn et al. presented a gesture recognition system based on electric fields produced by devices within the user's home environment [CGL*12, CMPT12]. In contrast to the methods discussed before, the human body acts as a receiver electrode while a wearable hardware component analyzes the received signal.

Intrabody Communication: Capacitive coupling can also enable two devices to exchange information using the human body as a communication medium [FL96, Zim96]. This principle was firstly investigated by Zimmerman for realizing Personal Area Networks [Zim96]. In his work, he points out that all communicating nodes require not only a capacitive coupling to the human body but also a coupling to a common ground potential. Since then, this principle was picked up by many researchers, for example to identify different users at a touch screen [DL01]. Park et al. proposed the Touch and Play (TAP) system [PKS*06] for creating a link between media storage devices and output devices through the human body. In a concrete application scenario, a photo was transferred from a digital camera to a printer.

Proximity and Grasp Sensing using Capacitive Coupling: In the domain of proximity sensing, capacitive coupling can be used for recognizing gestures [GPBKK13, SGB99], sensing user locations [STS*13] or their activities [GPBB*13, WKBS07a]. Also, capacitive sensing was used in wearable devices to measure muscle contractions for activity recognition [CAL10]. The approaches rarely transmit information, they rather use traditional capacitive sensing techniques. However, modulation approaches like code-division multiple access or frequency-division multiple access can also be found in this domain [GPBB*13, SGB99]. In order to discern different ways a user touches an object, capacitive sensing technology such as swept-frequency capacitive sensing [HSP12, SPH12] or time-domain reflectometry [WB11] were employed.

4.3.2. Competing Communication Technologies

Figure 4.13 shows three different methods for communicating with electric and magnetic fields. These methods can also be linked, for example, some types of RFID communications use a magnetic field for power transfer and an RF signal for information exchange. In the following, I will compare CapNFC to radio-frequency and inductive communication methods.

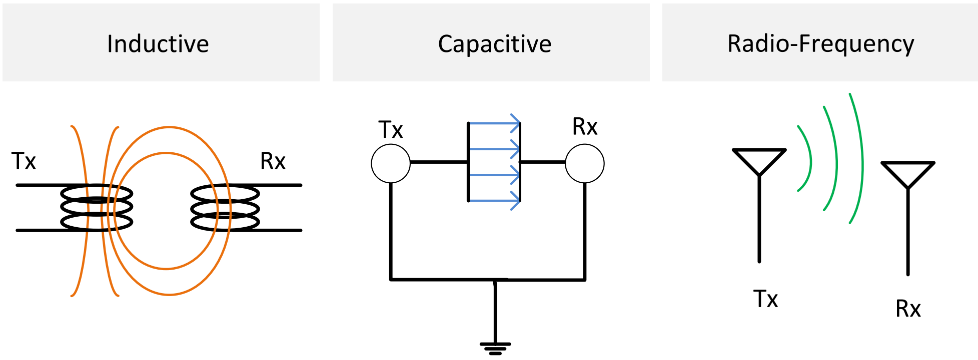

Figure 4.13.: Communication methods relying on electric and magnetic fields.

4.3.2.1. Long-Range Wireless Communication

The most important difference between CapNFC and long-range wireless communications can be seen in the supported transmission range and energy consumption. Technologies like Bluetooth or ZigBee support transmission ranges of 1 - 100 m, whereas CapNFC usually operates in a range up to 15 cm. Compared to wireless communications, CapNFC is easier and cheaper to implement, and consumes significantly less energy while transmitting and receiving data. On the other hand, using wireless communications offers several advantages like a full-duplex communication stack, greater distances and higher data rates. Therefore, the desired communication range, constraints in energy consumption and communication abilities are the most important decision factors.

4.3.2.2. Inductive RFID / NFC

In contrast to Capacitive NFC, Inductive NFC (RFID) has found many applications due to its commercial breakthrough and available passive tags. RFID was used in smart environments for interacting with everyday objects, for example coupling a mobile phone to a stereo device [CPL12] or enabling an information flow between touch-enabled surfaces [FWK*13]. Moreover, it is used as a bridging technology for establishing a confidential pairing with other wireless communication services and for mobile micropayment solutions [DL12]. In the following, we compare the power consumption of our CapNFC implementation to an ultra-low power RFID (NFC) implementation given in [Tex14b]. Prior to this discussion it must be noted that both implementations have very different levels of technological maturity and RFID's communication concept provides more elaborate features [NFC14]. RFID relies on a reader-central infrastructure supporting passive tags, whereas CapNFC's communication mechanism is based on broadcasting and a decentralized infrastructure. Also, the data rates strongly differ between both exemplary implementations (CapNFC: 2kBit/s; RFID/NFC: 424kBit/s) [Tex14b].

In order to compare the two approaches, we apply a simplified measure of energy consumption per bit $E_{bt} = \frac{\text{Energy Consumption}}{\text{Maximum Data Rate}}$. An overview of CapNFC's energy consumption in comparison to other communication methods is depicted in Figure 4.14. When transmitting, our CapNFC implementation's energy consumption is lower (360 nWs/Bit) than the given RFID implementation (825 nWs/Bit) [Tex14c]. While the RFID reader requires more than twice the energy per bit, it is also able to power passive tags, which CapNFC does not do. A comparison of both communication channels, i.e. ignoring the peripherals' energy consumptions, shows a larger difference. An output energy of 235 nWs/Bit is typical for RFID [STM14,Tex14c]. CapNFC's output energy varies from 0.225 nWs/Bit while transmitting through

	Energy Consumption incl. Peripherals	Output Power	Distance
CapNFC	360 nWs / Bit (TX) 7500 nWs / Bit (RX)	< 2.475 nWs / Bit	< 0,2 m
RFID/NFC	825 nWs / Bit (1)	235 nWs / Bit (1,2)	< 0,2 m
ZigBee	288 nWs / Bit (3)	4 nWs / Bit (3) (before antenna)	< 100 m

1. STMicroelectronics Inc. STRFNFCA Near field communication transceiver, 2014.
2. Texas Instruments Inc. TRF7970A Multiprotocol Fully Integrated 13.56-MHz RFID and NFC Transceiver IC, 2014.
3. Texas Instruments Inc. CC2520 DATASHEET 2.4 GHZ IEEE 802.15.4/ZIGBEE® RF TRANSCEIVER, 2014
4. Micro Crystal Switzerland. OV-7604-C7 32.768 KHz Oscillator, 2014

Figure 4.14.: CapNFC's energy consumption compared to other communication technologies.

air to 2.475 nWs/Bit while transmitting through the human body. This suggests that more mature CapNFC implementations can lead to an energy consumption which is many magnitudes smaller than RFID. The main reason is the small amount of energy needed to charge the capacitance between transmit electrode and environment. CapNFC's energy consumption while receiving information is mainly induced by the transimpedance amplifier and the signal processing on the microcontroller. As we did not optimize the transceiver board for low-power operation, signal processing leads to an energy consumption of 7.5 μWs/Bit. Nevertheless, we expect a great potential for reducing energy consumption in the receiver part.

Due to CapNFC's early development stage, only few approaches have been made to create passively powered communication devices. However, power transmission using capacitive coupling, as employed in [Mur11], proves the potential of passively powered capacitive sensor tags. Even though the comparison of energy consumption between both technologies is very difficult, especially the low energy consumption for transmitting information features CapNFC for use in ultra-low power devices with unidirectional communication abilities. We also see potentials in the combination of energy harvesting techniques [PS05] to realize passive tags.

Inductive RFID / NFC		**CapNFC**	
+	Tagging of objects is very simple	-	Flexibility and operating modes introduce complexity
-	Large objects are not easy to tag	+	Very flexible for different object sizes using conductive parts as electrodes
-	Limited interaction modes	+	Proximity & indirect touch sensing, intrabody communication
+	Passive tags are available	o	Passive tags can be realized
-	High peak currents during communication	+	Very energy-efficient information exchange
+	No ground-coupling is required	-	At least two electrodes needed

Figure 4.15.: Comparison of RFID with CapNFC properties for UbiComp.

RFID does not require a common ground between communicating nodes which is desirable for many applications However, the need for establishing a capacitive coupling to a common ground in CapNFC enables numerous opportunities for embedding interactive properties in an interaction system, as described in the technical framework section. The different CapNFC operating modes allow for intrabody communication, proximity awareness and indirect touch sensing. Comparing CapNFC's electrode placements to the placements of RFID antennas, one can achieve similar form factors, as described in the technical evaluation and the following case studies. The need for a common ground in CapNFC requires an additional ground electrode, raising the complexity of tagging an object. A very interesting property in CapNFC is the equipment of large objects with communication abilities, which would require multiple RFID readers. For example, a bed can be equipped with simple conductive wires to build a large receiver area. In conclusion, the decision of using CapNFC or RFID strongly depends on the application scenario, an overview of our discussion is outlined in Figure 4.15.

4.4. Case Studies: CapNFC in Ubiquitous Computing

Sensing object manipulations in combination with Capacitive Near-Field Communication offers an elegant solution for perceiving a user's environment as well as interesting ubiquitous interaction opportunities. CapNFC has a number of useful properties that are highly desirable in Ubiquitous Computing:

1. Instantaneous, infrastructure-free communication: Using CapNFC, interaction is independent from a server-based infrastructure or central access points. Instead, a multitude of devices within range communicates instantaneously based on a stateless communication protocol.
2. Natural interaction: A user can directly identify and combine the objects required for interaction. By employing additional sensors within an object, the knowledge of object manipulations and movements can be exploited and published in a short-range context. Combining the different CapNFC operating modes can be utilized for enriching the interaction, for example by detecting touch, proximity or communicating messages through the human body.
3. Low energy consumption: Capacitive Near-Field Communication requires significantly less energy than high-frequency wireless communication.
4. Low hardware requirements: Transmitting information is as easy as toggling a pin on a microcontroller, without requiring any additional integrated circuits. Moreover, the layout of transmit and receive electrodes is very flexible in terms of materials, shapes and sizes.

In the following, we present three case studies that benefit from CapNFC and combine the different operating modes in a useful way. In the first use-case, we present a real-world example where blind users interact with a PC using tangible objects. The second use-case is a conceptual study of natural interaction with an everyday object. Thirdly, we present an activity recognition system for sleeping analysis.

4.4.1. Cast Study 1: Tangible Interaction for the Blind

In this case study, we present a prototypical implementation of a tangible interaction system for visually impaired users. Especially for this target group, tangible objects provide an easy and intuitive way of symbolically accessing computer functions [MRB10, RHR10]. To develop a feasible and ergonomic solution, we work in close cooperation with a company specialized in developing products for the blind.

In order to use a computer efficiently, blind people are obliged to learn a high number of keyboard shortcuts and commands that are used to automate the computer. Symbolic access methods that link tangible objects to computer functions can help those people to execute a very specific set of actions, such as opening the weather report or reading their emails. Especially when considering the usage scenario of controlling continuous system properties, such as the reading speed of a screen reader, the usage of tangible objects employing CapNFC enables a very fast and natural way of interaction. Figure 4.16 shows different tangible objects that were realized for this use-case. A transceiver is connected

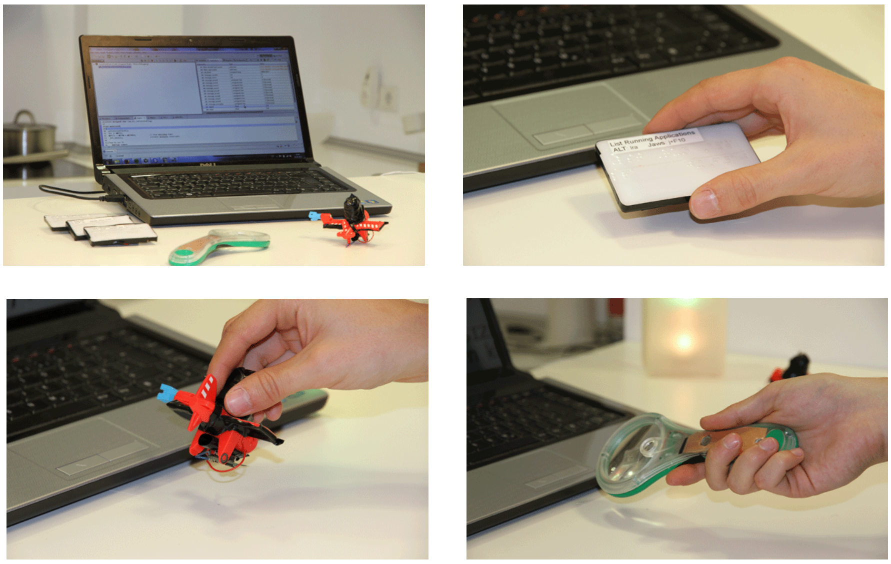

Figure 4.16.: An exemplary workplace equipped with CapNFC-enabled objects for symbolic access to a very specific set of computer functionality. The transceiver is connected to the PC, having a receiver electrode placed under the user's desk.

to a PC and the receiver electrode is deployed under the desk. The first object we realized is a plastic card with braille text. Each card represents a command that is carried out by the computer, for example 'list running applications'. Therefore, when the card is being touched (*ground-coupling by touch*), the card's ID and its acceleration data can be sent to the computer. When the card is successfully identified, a double tap on the card executes the corresponding command. The second object is a toy airplane used for regulating the reading speed of the screen reader. Leaning the plane forward and backward increases and decreases the speed. The third object - a magnifying glass - operates in two modes: touch and proximity. When the user tries to find the magnifying glass on the desk, the computer is able to play a sound when the hand moves over it. When the glass is touched, it can be leaned in two directions to control the position of the screen magnifier. By employing the signal-to-noise ratio at the receiver, the magnification level can be increased and decreased depending on the glass's distance.

As the behavior of each object is very specific, we implemented context-sensitive audio hints for each object by shaking it. Moreover, acoustic notifications indicate when an object was successfully recognized or when the object is put away. Our interviews and initial experiments with three blind users revealed that the setup must be chosen very carefully, according to the target person or group. As the initial learning process of keyboard-based systems is very difficult, elderly, children, and persons with multiple disabilities represent a suitable target group for the system. People who are solely visually impaired and have no difficulties in learning are probably not going to benefit from the system, as a high productivity can only be achieved when using keyboard shortcuts. Nevertheless, in working life, computer workplaces for blind persons are often equipped with highly specific hardware, in which certain tangible objects could complement traditional interaction methods. The long battery life and very low cost of CapNFC-enabled tangible objects represent very important factors for the acceptance of the tangible interaction system.

Besides the interviews, we evaluated different objects involved in the case study to demonstrate the suitability of CapNFC communication. Therefore, we observed the SNR as well as the bit-error-rate before data correction (BER $= \frac{\text{\# incorrect bits}}{\text{total \# of bits}}$). Our measurement results are listed in Figure 4.17, showing

Operating Mode	Receiver	Sender	SNR [dB]	BER
Ground-coupling by touch	Laptop (electrode in the desk)	Braille Card (object distance 5 cm)	10.47	2.50%
		Airplane (object distance 5 cm)	11.40	1.40%
Ground-coupling by proximity	Laptop (electrode in the desk)	Magnifying Glass (hand distance 1 cm)	8.56	4.17%
		Magnifying Glass (hand distance 2 cm)	7.93	4.39%
		Magnifying Glass (hand distance 5 cm)	6.82	7.40%

Figure 4.17.: Case study 1 - signal-to-noise ratio and bit error rate evaluation for different tangible objects above a table.

that communication for *ground-coupling by touch* is very reliable for object distances of 5 cm. As both evaluated transmitting objects have rather small electrodes reaching from 1 cm^2 to 2 cm^2, the SNR drops to less than 3 dB at distances above 8 cm. *Ground-coupling by proximity* easily allows for recognizing when a hand approaches the magnifying glass. Therefore, the object's ground electrode was chosen to be 3 cm^2. As depicted in the table, it enables a developer to reliably detect an approaching hand at distances of 5 cm. Due to the bigger ground electrode, the magnifying glass is sufficiently grounded by the table to retain communication abilities without human presence.

4.4.2. Case Study 2: Interaction with Everyday Objects

As computing and communication technologies advance, more and more household devices and tools gain 'intelligence' and extended configuration possibilities. For example, lamps with adjustable light color and/or time-dependent lighting patterns have been commercially available for some time. However, such advanced features also require more powerful user interfaces in order to leverage their capabilities. Many users usually suffer from poorly designed interfaces, which induced us to transfer a natural interaction technique on a conceptual lighting control.

Figure 4.18 depicts a way users may naturally interact with the lamp. Therefore, different objects were equipped with CapNFC tags, whereas the lamp employs a transceiver. We stored CapNFC tags in the bottle caps of three bottles filled with differently colored water. When the user moves a bottle cap into the proximity of the lamp, the tag's acceleration values can be received by the lamp. When the bottle is leaned in the way someone would do to pour out water, the lamp starts filling up in the color of the bottle's contents. The ground-coupling is achieved by a small conductive silver wire, which is touched by the user's hand. A rubber is used to erase the lamp's content, making use of the internal accelerometer. It also features *ground-coupling by proximity* when being placed on the lamp's surface. In this case, approaching the rubber in distances up to 10 cm turns the lamp either on and off. Moreover, objects like a lighter or a whale are able to trigger specific lighting profiles like fire or ocean animations. We also placed three transmit electrodes, representing red, green and yellow, on the table. The three tags were connected to the common ground, enabling them to transmit messages through *intrabody communication*. As soon as the user touches both the lamp and the colored region, the lamp will fill up with the selected color. Using a mobile phone with a transceiver on its back allows the user to directly transmit color values to the lamp using a color picker. In this setup, the device is automatically grounded when touching its metal case.

Figure 4.19 shows an evaluation of the smart objects involved in the case study. Despite the small transmitting electrode, the lighter showed a very good performance enabling communications in dis-

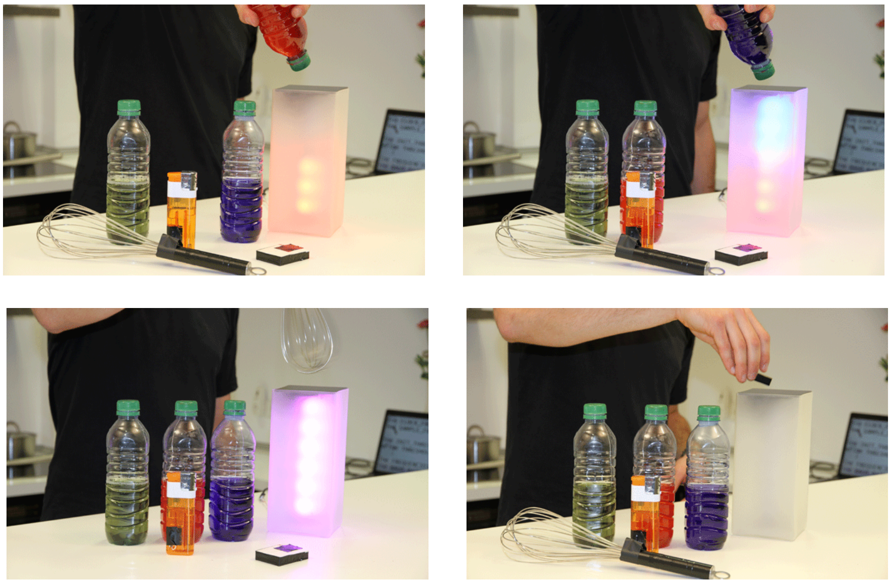

Figure 4.18.: An exemplary workflow for interacting with a smart lamp: The bottles are used to virtually fill up the lamp with the corresponding color. Moving the whisk above the lamp will mix the colors, whereas a gesture with the rubber gum switches off the lamp.

Operating Mode	**Receiver**	**Sender**	**SNR [dB]**	**BER**
Ground-coupling by touch	Lamp	Bottle (object distance 5 cm)	9.36	4.14%
		Lighter (object distance 5 cm)	10.06	1.91%
Intrabody with common ground	Lamp	Colored Area (coupling with 1 finger)	9.69	2.33%
		Colored Area (coupling with 2 fingers)	9.92	2.95%
		Colored Area (coupling with 1 hand)	10.50	1.60%

Figure 4.19.: Case study 2 - evaluation of communication properties for interacting with everyday objects.

tances up to 10 cm. Here, the lighter's conductive cap was used as a transmit electrode, which even works when the flame is enabled. In the future, we plan to exploit the increasing signal level when using *intrabody communication* to adjust the lamp's brightness by adding and removing fingers on its surface. The individual signal levels make it difficult to classify such gestures, as the signal has varying offsets influenced by skin conductivity and differences in the contact area. To conclude, this case study underlines the possibility of implementing new interaction paradigms with CapNFC by combining arbitrary objects to realize a natural way of interacting with everyday objects.

4.4.3. Case Study 3: Activity Recognition and Wearables

Activity recognition in combination with wearable and stationary sensors has been an emerging research topic for the past decade. Various approaches led to a wide variety of sensor-based approaches, for example to recognize activities of daily living [ASLT05, LJA*12, SHVLS08a] or studying various types of sleeping behavior [BV12, SSvL12].

Figure 4.20.: A bed that is able to receive messages from multiple body-worn sensors, for example a wrist-worn accelerometer. The accelerometer is grounded using ground-coupling by touch (inner electrode) and transmits its sensor values to the bed (outer electrode).

Transmitting information from multiple sensors to a smart bed can be used for online activity recognition or transferring sensor data collected during the day. Communication distances from the human body to a bedcover are very short, and are therefore well-suited for using CapNFC. In order to demonstrate CapNFC's potentials for sleep recognition, we built a smart bed in combination with a wearable sensor as depicted in Figure 4.20. The bed incorporates a transceiver with long electrodes placed under the bedcover. The wrist-worn accelerometer periodically reports its sensor data to the bed. It applies *ground-coupling by touch* with ground electrodes placed near the user's skin and a transmit electrode placed on the wristband's outer part. Furthermore, different regions in the center area near the thigh were equipped with stationary electrodes, transmitting information through *intrabody communication* when the person lies on them for identifying active regions. This arrangement exemplifies one of CapNFC's strengths: The supported flexible electrode layouts allow for easily equipping large objects with communicational abilities. Due to the different operating modes it is possible to join both, stationary and wearable sensing.

Operating Mode	**Receiver**	**Sender**	**SNR [dB]**	**BER**
Ground-coupling by touch	Bed (copper)	Wristband (object distance 1 cm)	20.31	1.12%
		Wristband (object distance 8 cm)	11.79	1.38%
Ground-coupling by touch	Bed (thread)	Wristband (object distance 1 cm)	17.65	1.52%
		Wristband (object distance 8 cm)	11.85	1.84%
Intrabody with common ground	Bed (copper)	Stationary transmitter (copper)	16.67	1.61%

Figure 4.21.: Case study 3 - communication performance for stationary and wearable objects communicating with a smart bed.

Figure 4.21 shows the properties of the communication link between the bed, the wristband, and the stationary electrodes. We evaluated two electrode materials within the bed, multiple conductive threads and tiny copper bands placed underneath the bed sheet. Despite the bigger area covered by the copper bands and their outstanding conductivity, the performance is very similar to using conductive threads. Conductive threads provide a lot of advantages in this case, they are very flexible and can be integrated easily into different kinds of fabric. The evaluation also shows that it is feasible to combine body-worn and stationary appliances to measure bedding postures and fine-grained physical parameters. Nevertheless, when considering wearable devices the data link may always be interrupted by posture changes. Also, using many highly active devices leads to a high number of message collisions.

4.5. Summary

As shown in the case studies, CapNFC has proved to be a very suitable technology for ubiquitous interaction and perception. Using CapNFC, it is possible to bridge the gap between a high number of smart objects, low cost, low power consumption and highly interactive system designs. Therefore, the technology represents a suitable companion to RFID with many exciting benefits in interaction design. The implementation is still in a very early stage, heading towards many types of possible applications. A combination of wearable and stationary sensors, as described in the last case study, represents one of the most interesting aspects for future research.

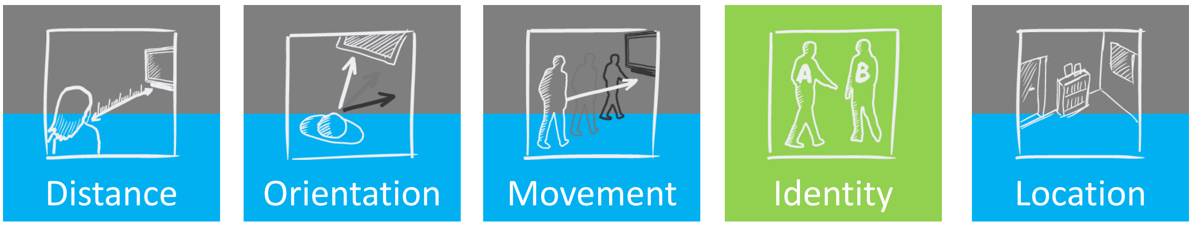

Figure 4.22.: Using CapNFC and OpenCapSense, I am able to cover all proxemic interactions presented in [GMB*11]. OpenCapSense covers the dimensions of distance, orientation, movement, and location, CapNFC covers identity.
[GMB*11] © 2011 Association for Computing Machinery, Inc. Reprinted by permission

Based on OpenCapSense and CapNFC, I am now able to provide means for measuring all proxemic interaction dimensions, as illustrated in Figure 4.22. Both contributions enable to create highly interactive system designs, that are unobtrusive and consume a low amount of energy. The two contributions can be seen as the technological foundation for the latter part of my thesis.

5. Object Recognition based on Capacitive Proximity Sensing

Having approached the basics of sensing human interactions based on capacitive coupling, I now take a closer look at making sense of data. This is required to finalize the perceptional capabilities in terms of proxemic interaction dimensions. As a singular capacitive sensor may only deliver a one-dimensional piece of information, fusing data of multiple sensors is necessary to obtain higher-level information. Moreover, when redundant information is sensed, sensor noise can be compensated. In recent years, many approaches on 2D-object recognition systems have been investigated and incorporated in commercial products, such as touchscreens. Moving from the 2-dimensional space to three dimensions, I experienced a lack of suitable methodology in capacitive object recognition. Here, the goal is to extract the 3D-configuration of an object, such as a human hand. These configurations can include multiple degrees-of-freedom, for example the pitch, yaw and role of a hand. Considering smart environments, object recognition may also be used to reconstruct whole-body parameters. Especially in this domain, many approaches for object and movement recognition rely on classification [SPH12, CMPT12]. In contrast to discrete inferences, recognizing continuous properties reveals significantly more information about human movements.

Due to the limited resolution of capacitive proximity sensors, it is necessary to make assumptions about the recognizable object in prior [SGB99]. This problem must be investigated very carefully as new objects might lead to confusions. Basic 3-dimensional object recognition is already available in commercial systems like Microchip GestIC [Mic14a]. It is able to extract a 3D-location in space based on four sensor values. Other systems, such as Cypress TrueTouch, are able to extract 2.5-dimensional hovering information in smartphone touchscreens [Cyp12]. However, there are currently only few generic approaches on object recognition in capacitive sensing. For these reasons, I contribute to the second research goal on data interpretation and fusing with an object recognition method. As depicted in Figure 5.1, the method covers the proxemic interaction dimensions of distance, movement, orientation, and location.

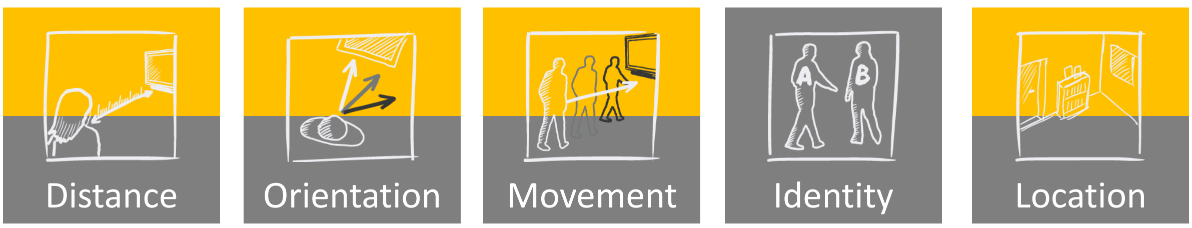

Figure 5.1.: Swiss-Cheese Extended provides means for data processing in the proxemic dimensions of distance, movement, orientation, and location.
[GMB*11] © 2011 Association for Computing Machinery, Inc. Reprinted by permission

The method's foundation is the *Swiss-Cheese Algorithm*, presented as future work by Smith et al. [Smi96]. Based on this briefly described algorithm, I formulated a method for object recognition with capacitive sensors in my master's thesis [GP12]. It enables recognizing and tracking multiple objects in real-time using shunt mode sensing. In this dissertation, I provide a revised and generalized version of the method, called *Swiss-Cheese Extended*. The first version of Swiss-Cheese Extended was published in [GPBKK13]. Later, Yannick Berghöfer's master thesis investigated extensions to the algorithm that were incorporated into the method. Besides the method's evaluation presented in my master's thesis, I present an evaluation on a novel use-case [Ber12]. A new forward model for loading mode measurements is investigated and new processing steps on preprocessing and gesture recognition are introduced. Figure 5.2 shows the two setups on which the method is evaluated.

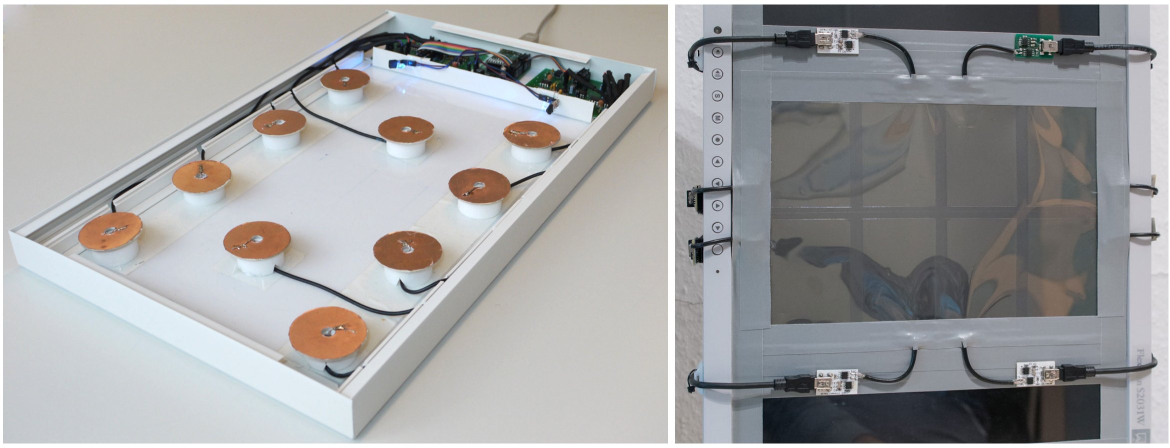

Figure 5.2.: The method is evaluated on two custom-built gesture-recognition systems. One applies shunt mode for free-space interaction [GP12] (left) and one loading mode for interacting in front of a display [Ber12] (right).

This remainder of this chapter is based on the papers [GPBKK13, GPB12] and my master's thesis [GP12]. The chapter contains findings from Yannick Berghöfer's master's thesis [Ber12], which are referenced in the corresponding sections. The use of 'we' in this chapter refers to the papers' authors Tobias Grosse-Puppendahl, Andreas Braun, Felix Kamieth, and Arjan Kuijper.

5.1. Swiss-Cheese Extended

Inferring different object parameters from sensor data is a complex task. An exact solution would require solving electric field equations for multiple objects and electrodes. This calculation is too time-consuming for real-time calculations in embedded systems and requires including numerous detailed environmental parameters. Another prevalent issue of capacitive proximity sensors is related to a certain ambiguity in sensor readings. Considering a single sensor and its generated electric field, a small object that is close to the sensor might result in the same reading as a larger object at an increased distance [Bax96]. Thus, a model is required that approximates the behaviour and influence of objects within an electric field. There are various practical solutions to build such a model. Typically the actual shape of the desired object is approximated by simple geometric shapes that are easier to process, e.g. spheres for modeling hands or cylinders for modeling arms [SGB99]. Reducing the complexity even further they are often considered uniform in size and shape which allows associating sensor values to a specific distance [BH11]. However, reducing the number of parameters reduces available information accordingly. When considering more complex scenarios, such as multi-hand gesture interaction, it becomes necessary to handle objects with multiple degrees-of-freedom and objects that are linked together with various geometric constraints. Therefore a method is required that considers these restrictions and allows recognizing and tracking the state of various objects in real-time.

The basic idea of the Swiss-Cheese-Algorithm is to detect objects using elimination [Smi96], as depicted in Figure 5.3. Initially it is assumed that the objects may be located at any position in the interaction space. Based on the measurements of each sensor we can make assumptions about the space in which no object may exist, reducing the probability of object presence around a certain proximity to the sensor. Combining the readings from many sensors we end up with a structure not unlike a Swiss cheese, with regions that may contain an object and others that are distinctively empty. While the basic idea of this algorithm has been outlined in the past as an outlook on future work, there has not yet been any concrete implementation or theoretical formulation [SGB99]. In the following, we present a conceptual and mathematical foundation of the Swiss-Cheese-Algorithm and various extensions that facilitate object recognition and tracking.

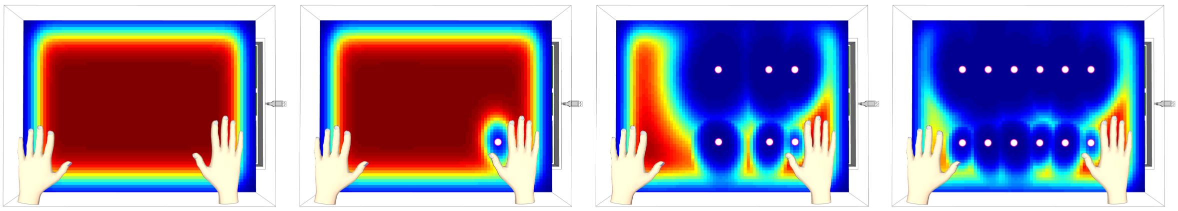

Figure 5.3.: Swiss-Cheese-Algorithm combining the knowledge of 0, 1, 6 and 12 sensors to recognize two hands. The figure shows a 2-dimensional layer of the Swiss-Cheese-Algorithm's outcome directly underneath both hands. White dots denote the center of an active sensor (receiver-transmitter combination). Red colors denote high probability of object presence (close to 1), while blue colors denote low probability of object presence (close to 0).

5.1.1. Method

In this subsection, we give a short overview about the processing steps of our object recognition and tracking method. The method is feasible for many different application scenarios and can be easily adapted. We illustrate these steps with our study of a multi-hand gesture recognition device, shown in Figure 5.4. We aim to determine the most likely configuration of body parts based on the readings of many distributed proximity sensors. One important requirement of the method is the applicability on environments where it is not feasible to deploy a large amount of sensors. Thus, we have to make preliminary considerations about the recognizable objects and their degrees of freedom. As a first step we define a volumetric model of the object to be recognized. Referring to our study of a multi-hand interaction device that is shown in Figure 5.4, we aim to recognize the 3D-positions and grabbing state of one or two hands. Therefore, the hands are modeled as boxes with a variable x/y-edge length and an (x,y,z)-position, resulting in a 5-dimensional descriptor, the *object state*. While the position of the center-of-gravity of this box is directly associated to the position of the hand, the edge length in two dimensions and their ratio to each other are used as indicator for the grabbing state.

In the algorithm's first execution step, a volumetric object is defined that encloses the whole interaction space, that we are calling cheese. This cheese can be regarded as a 3-dimensional pseudo probability distribution for object presence in each point [Smi96]. At the beginning of the algorithm, the presence of body parts is considered with equal probability everywhere, comparable to a cheese without holes. The algorithm has to cope with a high degree of ambiguity as sensors might deliver the same sensor reading for varying object sizes and distances. Thus, the algorithm can only make considerations about

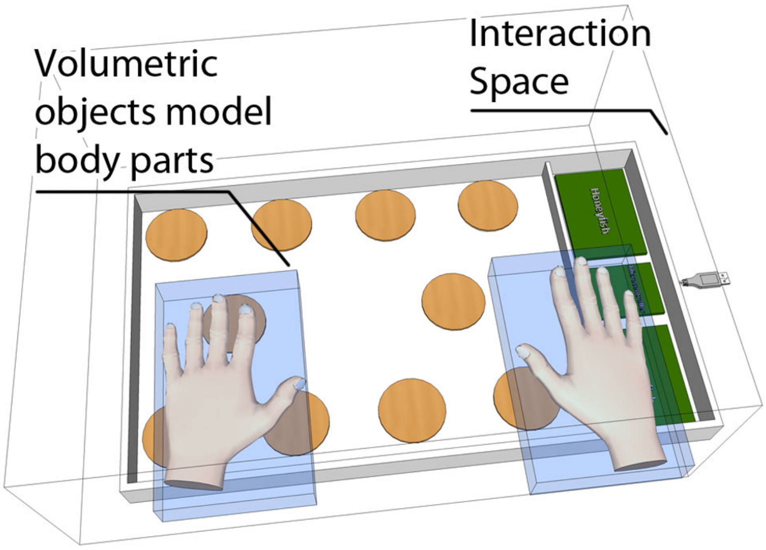

Figure 5.4.: Our multi-hand interaction device with hands modeled as volumetric objects. 10 copper plates are used as electrodes that build up an electric field to the user's hands.

the space around a sensor in which definitely no object is present. This space can be modeled using an ellipsoid around the sensor's center. In the following, these ellipsoids are cut out of the cheese for each sensor. The parts left over contain the objects we want to recognize.

Let us illustrate this procedure with an example shown in Figure 5.3. It shows a 2D-layer of the cheese that is located in the interaction space above the multi-hand interaction device presented in our study. Ellipsoids are cut out of that cheese and the position of the two hands is subsequently revealed. Afterwards, the defined volumetric models of the objects are fit into the remaining cheese to obtain a probability measure for different object part configurations. However, it is typical that a large portion of cheese remains, as the sensors are usually not able to constrain the interaction sufficiently in all directions. Thus, we associate a higher weight to the object state, when it is located closer to a sensor compared to states that are located at greater distances. Using the example of the gesture recognition device it is easy to see that if a hand is located 10 cm above the sensing plane, the probable object configurations are recognized at this distance and above.

In order to determine the most likely system states in real-time, it is not feasible to evaluate all possible object configurations. Especially when the number of targets increases or the object state vector's dimensionality is high, a systematic approach for finding the most probable object configuration has to be considered. Particle filters, also known as Sequential Monte-Carlo method, provide a solution to this problem. These filters can be incorporated to evaluate only the most probable object configurations based on a spatio-temporal relationship. Multiple objects can be tracked using separate instantiations of a particle filter, which enables us to track two hands in real-time with our multi-hand interaction device.

5.1.2. Object Recognition

The process of recognizing objects is outline in Figure 5.5. It starts with a set of calibrated and normalized sensor measurements. Additional filters can be applied on the raw sensor measurements, which should be chosen individually according to the use-case. For each sensor, a forward model exists that can predict the sensor's value for a unit absorber located at a given point p_i in the interaction space. When

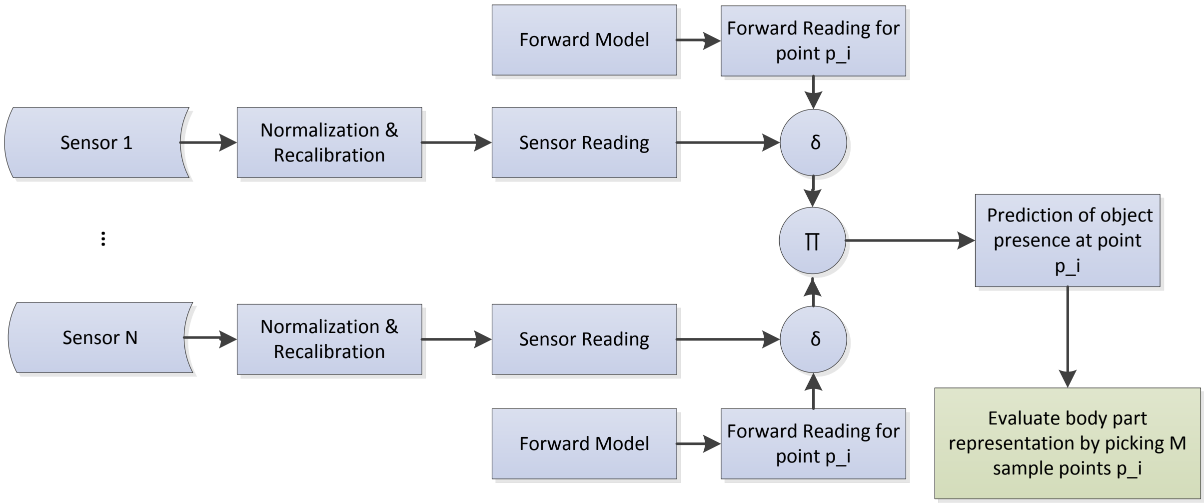

Figure 5.5.: The process of object recognition is based on a set of sensor measurements that are compared with forward readings.

comparing the actual sensor values s with the predicted forward readings f, it is possible to evaluate object presence at any point in space. Then, these points are used to evaluate a hypothetical body part configuration.

5.1.2.1. Normalization

The first step in signal processing is the normalization of raw sensor readings. The signal is mapped to a range between 0 and 1, while it is cut off below and above these bounds. Normalization makes it easier to cope with possible signal offsets and allow for recalibrating a sensor when environmental changes occur.

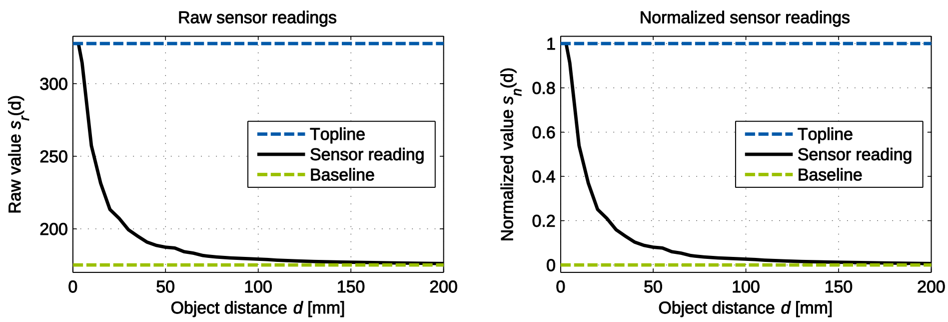

Figure 5.6.: Normalization depending on different object distances for loading mode measurements [Ber12]

Figure 5.6 depicts an example of normalizing sensor values s_n obtained from a loading-mode measurement [Ber12]. The baseline represents a system state when no hand is in proximity. In this case, one only measures the undesired parasitic capacitance which is induced by electronics and the environment. The topline is obtained when a hand is in the closest possible proximity to the sensor. This property is dynamically adjusted to the baseline and the maximum differential sensor response.

When the normalized signal drops below the current parasitic capacitance level, the baseline can be automatically recalibrated. This is usually the case when conducting objects are moved away from the sensors or when the sensors are placed on a less conductive surface than before. Recalibration becomes more difficult when parasitic capacitance is added instead of reduced. This can be due to temperature changes that induce a small drift, or new static objects placed besides the sensor. Such problems can be compensated slowly, for example by applying a small shift to the baseline over time. Alternative methods analyze the variance in the sensor readings, which changes when interacting with the sensor.

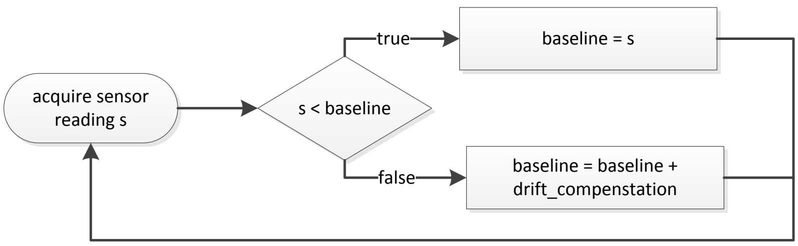

Figure 5.7.: Normalization procedure: When the baseline falls below a level, an instaneous recalibration is performed. Otherwise, drift compensation is added to the baseline.

Figure 5.7 depicts an exemplary process of normalizing sensor values for loading mode measurements. When the sensor value falls below the current baseline, it can be concluded that parasitic capacitance has been reduced. Therefore, a direct recalibration can be performed. Instead, when the sensor value is above the baseline, a small factor is added to the baseline for compensating drifts. This leads to a slow disappearance of objects which are not moved.

5.1.2.2. Forward Reading Model

The forward reading model is used to reconstruct a sensor reading f that one would obtain when a unit absorber is placed at a location (x,y,z). A unit absorber is a very small conductor and represents the nearest possible object in the environment of a sensor for a given sensor reading. However, it is possible that bigger objects cause the same sensor reading at greater distances from the sensor. The forward reading f can be seen as a prediction of a sensor reading for an imaginary unit absorber. Later, this prediction is used for comparison with the actual sensor value.

The *iso-signal shell* is a surface around a sensor on which a unit absorber causes the same sensor reading [SGB99]. Thus, this surface marks a volume around a sensor in which no object may be present. Due to the different electrode layouts, a sensor's reading depends on the direction in which an object approaches. For example, a unit absorber that is placed vertically above a sensor at a distance of $10cm$ could produce the same sensor reading as a unit absorber placed horizontally aside a sensor at a distance of $15cm$. These axis-dependent characteristics can be expressed by modeling the iso-signal shell as an ellipsoid, as shown in Figure 5.8. This ellipsoid is composed of three independent semi-principal axis $r_{x,y,z}$. When a unit absorber is placed at a point (x,y,z), it is located on the iso-signal shell if the following ellipsoidal condition is fulfilled:

$$\frac{x^2}{r_x^2}+\frac{y^2}{r_y^2}+\frac{z^2}{r_z^2}-1=0 \tag{5.1}$$

Depending on the applied measurement mode, a model function with different parameters for each axis is required to calculate the forward reading. This model function relates the forward reading to the

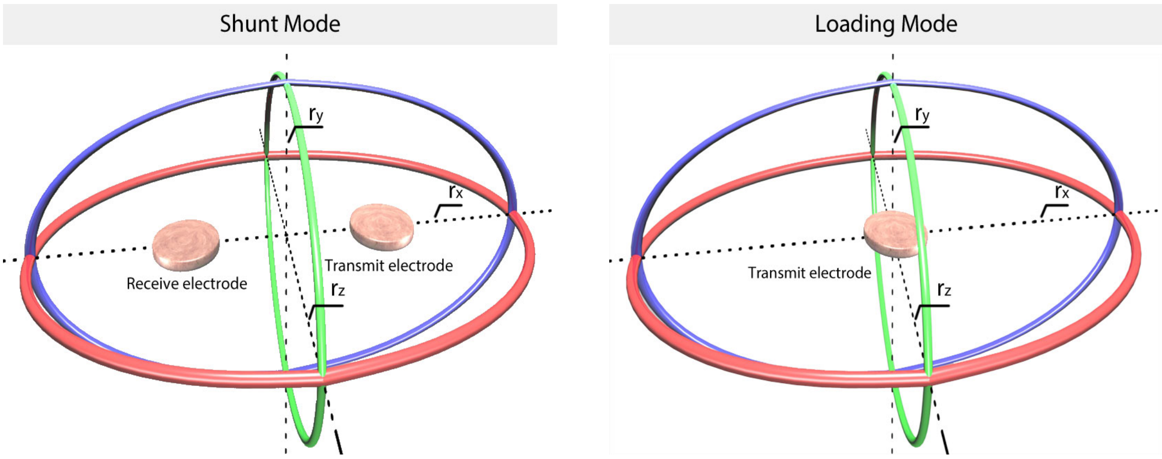

Figure 5.8.: An ellipsoid with three independent semi-principal axes $r_{x,y,z}$ models the distance of a unit absorber to the sensor's center. For shunt mode, the sensor's center is in the middle of transmit and receive electrode (left). The center of a loading mode sensor corresponds to the transmit electrode's location (right).

distance of a unit absorber. A simplified model can be based on the electric field strength in a given object location [SGB99] or the plate capacitor model.

Forward Reading Model for Shunt-Mode Measurements: Considering a dipole approximation based on point charges means that the electric field strength around a sensor decreases with the factor d^3, related to the distance d of a unit absorber from the sensor. Using this approximation we can derive the following model function for each axis $i = x, y, z$ to model the axis-dependent directivity: [Smi96]

$$f = 1 - \frac{1}{(\alpha_i + \beta_i r_i)^3}; i = \{x, y, z\} \tag{5.2}$$

$$\Leftrightarrow r_i = \frac{\beta_i}{\sqrt[3]{1-f} - \alpha_i} \tag{5.3}$$

The equation is composed of two fit parameters for a single axis, $\alpha_{x,y,z}$ and $\beta_{x,y,z}$. These fit parameters are applied to model the gradient of the electric field strength along the given axis. The fit parameters can be determined experimentally by moving a unit absorber along an axis and recording the sensor value in relation to the unit absorber's distance to the sensor. In the following step, the fit parameters are calculated using least-squares fitting.

We now combine the given fit function with the ellipsoidal condition to determine the forward reading for a unit absorber located at a certain point (x, y, z) on the iso-signal shell. Since the three semi-principal axis $r_{x,y,z}$ are unknown, we can replace them with the determined fit function given in Equation 5.3. We yield an equation with just one unknown variable, the forward reading f:

$$\left(\frac{x}{\frac{\beta_x}{\sqrt[3]{1-f}-\alpha_x}}\right)^2 + \left(\frac{y}{\frac{\beta_y}{\sqrt[3]{1-f}-\alpha_y}}\right)^2 + \left(\frac{z}{\frac{\beta_z}{\sqrt[3]{1-f}-\alpha_z}}\right)^2 - 1 = 0 \tag{5.4}$$

The equation cannot be resolved analytically to f, the result would be a polynomial of order 6. Thus, we choose to solve Equation 5.4 by minimization using the variable forward reading f. As a normalized sensor reading is restricted to a range of $[0, 1]$, the value of f can be determined efficiently with methods like the Brent algorithm [Bre02] in real-time. As an outcome of the minimization approach we can now determine a forward reading f each sensor would produce when a unit absorber is located at a given point.

Forward Reading Model for Loading-Mode Measurements: In his master's thesis, Yannick Berghöfer introduced a forward model for loading-mode measurements [Ber12]. Determining a forward reading f for this mode is not as straight-forward as for shunt-mode measurements. In particular, the z-axis has a different directivity than the x- and y-axis. This circumstance is described in more detail in the upcoming study section and is depicted in Figure 5.18.

The response on the z-axis is very similar to the one provided by the plate capacitor model, as described in 2.2.7.2. At closer hand distance larger residuals remain, which make it necessary to introduce four fit parameters $\alpha_z, \beta_z, \delta_z, \gamma_z$.

$$f = \frac{\alpha_z}{(r_z - \delta_z)^{\gamma_z}} + \beta_z \tag{5.5}$$

Modelling the x- and y-axis would require more complex electric field models. In this case, a mathematical interpolation is chosen to model the sensor response. Very similarly to the previous fit function, it utilizes another four parameters for fitting the sensor response: $\alpha_{x,y}, \beta_{x,y}, \delta_{x,y}, \gamma_{x,y}$.

$$f = \alpha_i \cdot e^{-\left(\frac{r_i - \beta_i}{\gamma_i}\right)^2} + \delta_i \quad ; i \in \{x, y\} \tag{5.6}$$

Inserting the two types of fit functions into the ellipsoidal condition yields the forward model. A solution can be approximated by minimizing that function as performed previously on the shunt-mode forward reading model. The forward reading f for a unit absorber located at point (x, y, z) can be obtained by the following equation.

$$\frac{x^2}{\left(\sqrt{-\ln\left(\frac{f}{\alpha_x} - \delta_x\right)} \cdot \gamma_x + \beta_x\right)^2} + \frac{y^2}{\left(\sqrt{-\ln\left(\frac{f}{\alpha_y} - \delta_y\right)} \cdot \gamma_y + \beta_y\right)^2} + \frac{z^2}{\left(\sqrt[\gamma_z]{\frac{\alpha_z}{f - \beta_z}} + \delta_z\right)^2} - 1 = 0 \tag{5.7}$$

Considering calculations on embedded systems, many forward models are too complex for real-time calculations. Therefore, it is more feasible to provide forward models that work with linear interpolation based on a number of previously determined support points.

5.1.2.3. Prediction of Object Presence

We can only make considerations about the space around a sensor in which no absorber can be located. This space is limited to the distance between a unit absorber and a sensor that can be regarded as the nearest possible object. In this step, spaces in which no object may be located are cut out of the pseudo probability distribution. To evaluate a point $\vec{p}$ in space, the forward reading f is subtracted from the actual sensor reading s that was measured, resulting in $\delta = f - s$.

To emphasize the meaning of the value δ, consider a unit absorber located $10cm$ above a sensor that would cause a normalized sensor reading $s = 0.5$. Applying the forward reading model for an imaginary unit absorber in a distance of $5cm$ above the sensor would yield a forward reading of $f = 0.2$. When the distance of the imaginary object comes closer to $10cm$, a forward reading of $f = 0.5$ would be determined. At greater distances, for example $15cm$, one would yield a forward reading of $f = 0.8$.

When computing the difference $\delta = f - s$ of the actual reading s and the forward reading f, we can conclude that no object can be present for $\delta < 0$ and an object can be present for $\delta \geq 0$. However, based on this assumption one can only conclude that the nearest possible object is located at 10cm above the sensor, but it is also possible that a bigger object is located at a distance of $11cm$ or beyond.

The difference value δ is an input argument to a sigmoidal function, with two parameters μ_n (displacement) and γ_n (steepness) where $n = 1, 2, ..., N$ denotes the sensor:

$$P_n(\vec{p}) = \frac{1}{1 + e^{(-\gamma_n \cdot (\delta - \mu_n))}} \tag{5.8}$$

The function expresses that the probability of an object being within the inner region of an iso-signal-shell of a unit absorber is close to zero. In the space outside the iso-signal-shell, the prediction is close to one. When the steepness γ_e converges to infinity, then the sigmoidal function can be regarded as a simple Heaviside function. Lower values for that parameter can alleviate the effect of noisy measurements by expressing a level of uncertainty. In the following, the knowledge gathered from all sensors is combined:

$$P(\vec{p}) = \prod_{i=1}^{N} P_n(\vec{p}) \tag{5.9}$$

Thus, when all sensors are sure that an object may be present at a given point, the function $P(\vec{p})$ will evaluate close to one. If the point is within a space in which one or multiple sensors do not consider an absorber, the function evaluates close to zero.

5.1.2.4. Body Part Representation

Based on the previous findings, we are able to obtain a measure for object presence in a single point. In the following, an approach for determining the state of an object is presented. As explained in the overview of this section, an object state can be embodied by a location and the properties of a volumetric object. However, it is possible to employ more complex geometrical models that are composed of many volumes and must be described by a higher number of parameters. It is necessary to find a suitable compromise between the object's shape and the accuracy of the model. It is not viable to apply a fine-grained arm model for the distinction of different fingers if the required information is not contained in the sensor data.

The most probable object state can be determined by maximizing the average volume integral over the pseudo probability distribution in each point that is enclosed by the volumetric model. The volume integral over the function $P(\vec{p})$ is solved using a Monte-Carlo integration. V denotes the object's volume, M the number of Monte-Carlo samples, and $\vec{p}_i$ a sampling point within the object:

$$\iiint_V P(\vec{p}) d\vec{p} \approx V \cdot \frac{1}{M} \cdot \sum_{i=1}^{M} P(\vec{p}_i) \tag{5.10}$$

To obtain the Monte-Carlo integral of a function, a uniformly distributed set of points within the volume must be determined. For each of these points, the prediction of object presence is computed. With an increasing number of points, the error between the actual integral and the Monte-Carlo integral is minimized. However, a high number of points leads to computationally higher cost.

In order to obtain meaningful results, the object state parameters must be limited in a way that restricts the object position to the interaction space and the object shape to feasible variants. As the interaction space is usually not restrained by sensors in all directions, the object configurations that are closer to sensors must be weighted higher as object configurations that are far away from the sensors. This linear weighting can be accomplished by calculating the distance to the nearest sensor or the distance to a sensing surface, when the sensors are located in a plane.

5.1.3. Object Tracking

In the previous section, a method for object recognition was introduced. Using the Swiss-Cheese-Algorithm, it is possible to obtain a measure of probability for object presence in each point in space. This is the basis for the recognition of objects that can be modeled by basic geometric shapes, such as a box. The goal of object tracking is to estimate a system state, employing a set of measurements in real-time. This estimation does not only depend on the current time-step, but also on the system state's evolution in time. In single-target tracking, a system state can be expressed by a single object state, for example by the position of hand modeled by a box. When more than one object is tracked, the system state is the combination of all distinct object states. A *system model* incorporates the change of a system in time, whereas the *measurement model* is utilized to evaluate the probability of a hypotheses [IB98].

In order to determine the most likely system states, it is not feasible to evaluate all possible system states in real-time. Especially when the number of targets increases or the object state vector's dimensionality is high, a systematic approach for finding the most probable system state has to be considered. Particle filters reveal their strengths in the possibility to track many hypothesis about an object state in a spatio-temporal relationship. The concept makes them robust against occlusion and clutter [IB98]. This robustness can be exploited in capacitive proximity sensing, as measurement noise and fast movements pose comparable challenges on the filter. Moreover, maintaining these spatio-temporal relationships can enhance the recognition rate when objects leave the interaction area for a short time.

Tracking multiple targets poses various challenges on particle filtering. Standard particle filters are not suited for tracking a varying number of targets. In particular, the samples quickly converge to a single target when more than one target is present [VDP03]. To overcome this limitation, many extensions to particle filters were proposed [KMA01, VDP03]. When the number of targets T is known in advance, it is possible to represent the system state as a joint set of object states [VDP03]. Problems arise when the object states become more complex and the number of targets increases. In this case, the dimensionality of the system state vector quickly becomes unhandy and the system performance decreases.

In order to avoid a rising complexity with an increasing system state dimensionality, targets can be tracked independently with separate instantiations of a particle filter. It is essential to protect samples that represent local maxima of the probability distribution from extinction. Milstein et al. present the idea of a clustered particle filter [MSW02], that inspired our multi-target tracking approach. In each step, the determined probability distribution is clustered for a variable number of targets. For each cluster, a fixed number of samples is selected for the next sampling stage. In each step, the number of targets is determined from the variance of object states within a corresponding cluster. When the variance is high, the number of targets is increased and the clustering process is repeated. When no good object states can be found within a cluster, the number of targets is decreased.

Furthermore, the task of tracking newly appearing targets is not considered in standard particle filtering [KMA01]. When particles track an existing target, a newly appearing target can only be recognized by particles that migrate from the existing target to the new one. In order to solve this problem, we determine an initialization density directly from the sensor readings. When a sensor yields a reading that indicates a nearby target, particles with expected object states are randomly initialized in the neighborhood of that sensor. In each initialization phase, a fixed number of particles is distributed over the state space.

Samples are represented by a set $s_t^{(n)}; n = 1,...,N$ with corresponding weights $\pi_t^{(n)}$, cumulative weights $c_t^{(n)}$ and a cluster $u_t^{(n)} = 1,...,C$. Each sample is connected to a single cluster, while the number C of possible clusters is estimated in each time step. In the following, a multi-target tracking algorithm is presented, as described in [GP12]:

1. For each cluster identified in the previous time-step, **select** a set of samples $s_t^{'(n)}$ by drawing $(M-L)/T$ samples from the set $s_{t-1}^{(n)}$. The samples are drawn with replacement, samples with high weights are drawn more often than samples with low weights.
2. **Initialize** L samples directly from the sensor readings. When a sensor outputs a low sensor reading, it is very probable that a target is present in its surrounding. The number of initialization samples $L_k, k = \{1,...,K\}$ per sensor with a reading $z_k \in [0,1]$ can be estimated by a weighted average:
$$L_k \approx \frac{1-z_k}{\sum_{i=1}^{K} 1-z_i} \cdot L \tag{5.11}$$
In the next step, the samples are randomly distributed in the region around a sensor. Object state parameters, that do not include the (x,y,z)-coordinate are set to a fixed value. It is noteworthy, that the generated samples are not related to a cluster yet.
3. **Predict** the object state at time t to retrieve a set of samples $s_t^{(n)}$. The system model incorporates the velocity of a particle in the previous two time-steps. This allows a fast tracking of movements:
$$s_t^{(n)} = 2 \cdot s_{t-1}^{(n)} - s_{t-2}^{(m)} + B \cdot w_t^{(n)} \tag{5.12}$$
A movement is diffused by a vector w_t, containing normal distributed random values, multiplied with a matrix B. This random part is utilized to incorporate acceleration and deceleration. In most cases, B is the identity matrix when the random parts do not depend on each other.
4. **Measure** the probability distribution for each sample and weight the particles according to the measurement result. In this step, the measurement model $p(\vec{s}|Z)$ is employed, with a succeeding normalization to meet the constraint $\sum \pi_t^n = 1$:
$$\pi_t^n = \frac{P(s_t^{(n)}|Z)}{\sum_{i=1}^{N} P(s_t^{(i)}|Z)} \tag{5.13}$$
5. **Cluster** the set of M samples by computing T new cluster centers and assigning a cluster to each sample. Clustering can be performed using the K-Means algorithm, that iteratively determines cluster centers by minimizing the distance of a sample to the nearest cluster center [Bis06]. When the object state variance within a cluster exceeds a threshold, the number of targets is increased and the clustering step is performed again. In case of a cluster having a very low average weight, or two clusters having a low distance, the number of targets is decreased and the clustering process is repeated.
6. **Resample** when the determined weights within a cluster have a high variance. Therefore, an estimate of the effective number of particles N_{eff} is computed for each cluster, as described by Arulampalam et al. [AMGC02]:
$$N_{eff} = \frac{1}{\sum_{i=1}^{N} \left(\pi_t^{(i)}\right)^2} \tag{5.14}$$

When this estimate for a cluster falls below a predefined threshold $N_{threshold}$, the set of samples is resampled. The resampling stage has no effect on other clusters, if they do not exceed the predefined threshold.

5.1.4. Target Management

The outcomes of the particle filtering approach for multi-target tracking are cluster centers that represent recognized targets and their properties. Target management is the task of keeping track of a uniquely identifiable target object through time and handling newly appearing and vanishing targets. Thus, the determined cluster centers that exceed an observation threshold are connected to a set of maintained target objects. A target object has a unique ID, a history vector of all identified cluster centers and threshold values that are used to compensate noisy detections and to maintain non-detected targets through measurement noise. For each time step, the target management assigns the recognized cluster centers to a set of maintained target objects. Therefore, the distances from the new cluster centers to the last assigned cluster center of all target objects is calculated. Then, the nearest cluster centers are assigned to the existing target objects. If more cluster centers than existing targets are recognized, the remaining unconnected cluster centers are used to create new target objects. When less cluster centers are recognized, targets that were not assigned to a cluster center are removed.

5.1.5. Interpolation

In general, the recognized cluster centers have relatively smooth trajectories over time. However, there might be variations due to noisy measurements and the probabilistic nature of particle filters. Thus, interpolation techniques can improve the continuity of trajectories in applications that track gestures. Moving average filtering of a target object's past cluster centers can smooth the trajectories and lead to higher precision. A moving average a for a target object with a history h of past cluster centers with window size L can be determined as follows:

$$a = \frac{1}{L}\sum_{i=0}^{L} h_{t-i} \tag{5.15}$$

An averaging approach with a fixed window size L faces the great disadvantage of increasing the latency to an unacceptable amount. When having smaller window sizes, the system's reaction time increases whereas the smoothness decreases. Most gesture-recognizing applications require low precision and latency while fast movements are performed. For tiny movements, the precision is considered to be more important than latency. Thus, we apply an adaptive moving average filter that determines the size of the input history cluster centers depending on the object's movement speed. Figure 5.9 shows an example of adaptive moving average filtering in a 2D-plane above a sensing device. In terms of spatial performance, the object trajectories are smoothed while edges are retained.

5.1.6. Gesture Recognition

Recognizing gestures addresses the problem of mapping a time series to a discrete class [Rab89, Cor01]. As this problem has been investigated in depth, I keep this section brief. In our case, the time series corresponds to a history of past cluster centers h. These cluster centers represent the multi-dimensional object state over time. First, it is necessary to segment the time series in order to determine a start and end point of the gesture. In the naïve view, the starting point of a gesture is the entrance of an object into the interaction space. However, this does not take objects remaining in the interaction space into

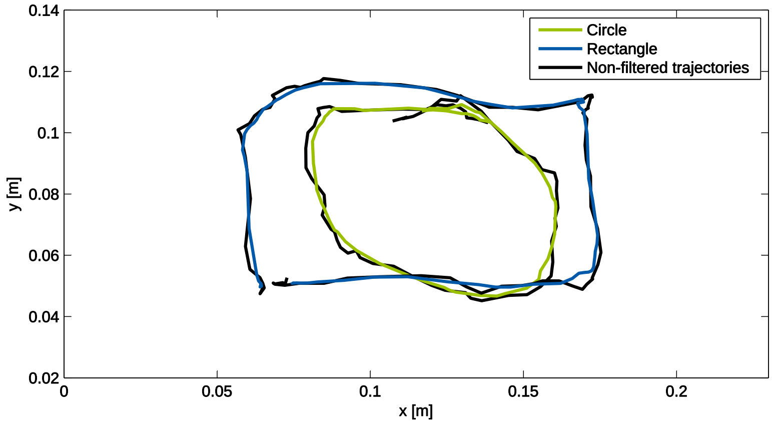

Figure 5.9.: 2-dimensional adaptive moving average filtering for object trajectories [Ber12].

account. Therefore, it is more feasible to set the start point of a gesture based on the basis of thresholding. The end point can be assumed after each time-step to keep the latency after performing a gesture low. Alternatively, it can be determined by stops or changes in movement.

Recognizing linear movements through space is very easy, as they can be recognized with heuristics. When the traveled distance of an object towards a certain direction exceeds a threshold, the gesture is recognized. Of course, this requires a well working segmentation beforehand. Other, more advanced approaches, are able to detect more complex gestures, such as circles or triangles. In his master's thesis, Yannick Berghöfer investigated the applicability of Hidden-Markov-Models (HMM) [Rab89] for gesture recognition [Ber12]. The chosen approach requires a discrete series of observations which is derived from the continuous time series. This set should be as small as possible, often requiring to reduce the dimensionality of recognized object states. This can be achieved by feature extraction, for example by extracting the movement angle between two succeeding 2D object positions, as shown in Figure 5.10.

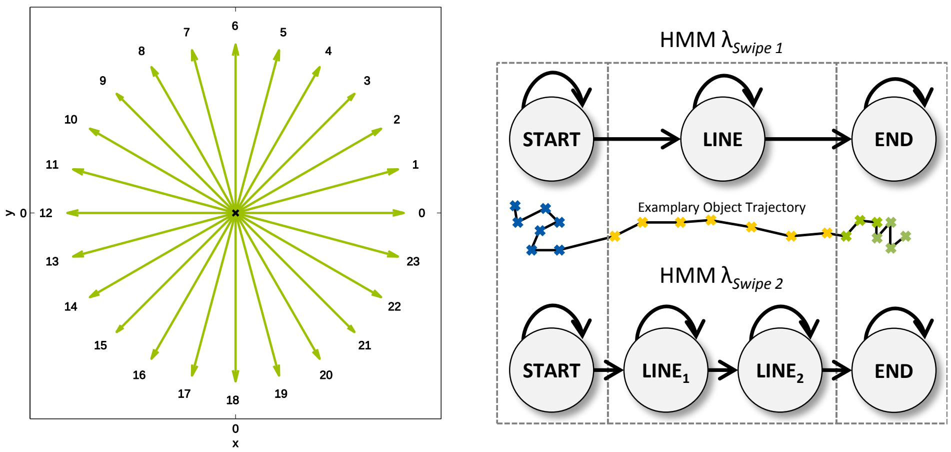

Figure 5.10.: Quantization of the object's movement angle to 24 discrete states (left). Yannick Berghöfer applied left-right-banded HMMs for recognizing gestures [Ber12] (right).

In the following, the discrete observation sequence is mapped to a set of HMMs, trained for each gesture. Berghöfer decided to use LRB (left-right-banded) HMMs that do not allow transitions to previous states. Moreover, it is not allowed to skip certain states, which is a difference to LR (left-right) HMMs. Figure 5.10 depicts an example of a LRB HMM. It maps the starting and ending parts of the object trajectory are mapped to the outer states. This models noisy detections at the beginning of each gesture. The inner states then represent the actual object trajectory.

An alternative method, which is easier to train, is Dynamic-Time-Warping (DTW). As part of a working paper, we investigate the use of this method for recognizing gestures directly on embedded systems [FGPK14]. DTW can be applied very easily on discretized observation sequences. The observation sequences are then compared to sequences contained template database. Mapping the two sequences to each other is performed in a non-linear fashion, since the scale between both sequences can be different. For example, this can caused by varying interaction speeds.

5.2. Study: Gesture Recognition Device

5.2.1. Prototype

We created a hardware platform for gesture recognition that is shown in Figure 5.2 [GPB12]. The platform operates in shunt mode and applies a combination of time-division and frequency-division multiplexing for parallel transmitter operation. The gesture recognition prototype uses two synchronized boards, each driving four transmitters and one receiver. The boards are able to receive transmitted signals from each other and can be easily extended. We transmit frequencies of 10-25KHz, sample the received signals at 100KHz and apply a Fast Fourier Transform for reconstruction of the transmitted signal amplitude. The transmitted sine-wave signals have a peak-to-peak amplitude of 5V. The multiplexing approach enables us to retrieve 50 samples per second for each receiver-transmitter combination.

The gesture recognition device consists of eight transmit electrodes located at both sides of the device and two receive electrodes that are placed in the center of the sensing plane. This mode of operation makes it possible to use 16 virtual sensors, one virtual sensor per receiver-transmitter combination. Applying this setup, we can detect fast multi-hand gestures above an area of 40 x 20cm with a maximum detection height of approximately 20cm.

5.2.2. Supported Gestures

Discrete gestures represent actions that trigger a discrete command, such as page turning. The recognition is based on the covered distance and movement direction of a target object, whereas the movement speed is of secondary interest. Thus, the movement history of each target object must be analyzed continuously. This processing takes place in each time-step applying a sliding window on the history of object states.

Swipe gestures, visualized in Figure 5.11, are well known in multi-touch applications and are often applied on image browsing or changing views [KFB*97]. This gesture type can be performed with a single target, but is also applicable on two targets moving in parallel. A swipe gesture is based on a movement parallel to a reference axis with low deviations to the orthogonal axis. Furthermore, it needs to be performed with a certain movement velocity. An average velocity in the x- and y-direction and the movement distance is calculated for a window with its past cluster centers. When the velocity in a single direction and the distance exceeds a threshold, a swiping action is recognized. In order to properly

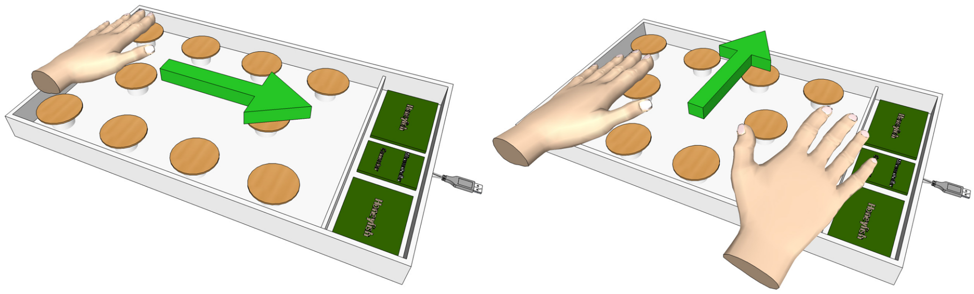

Figure 5.11.: Swipe gestures from left to right with a single hand and from bottom to top with two hands [GP12].

recognize a swiping action, the average velocity to the orthogonal direction must lie within an error threshold.

Continuous zoom and rotation gestures with two hands are shown in Figure 5.12 [BM86]. They are analogous to pinch/zoom and rotate gestures known from multi-touch applications [KFB*97]. In contrast to multi-touch, gestures are not performed using two fingers but with two hands. As soon as two hands are recognized, the corresponding angle between the hands with respect to the device's longitudinal axis is calculated, representing the desired rotation. The distance between both hands is mapped to a zoom factor.

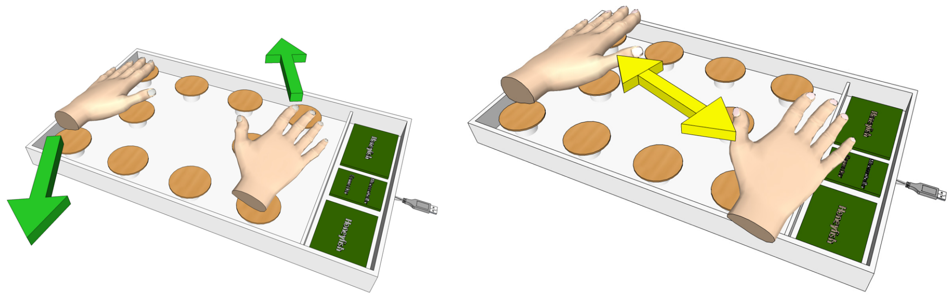

Figure 5.12.: Combined zoom and rotation gesture (green) and the corresponding zoom and rotation axis (yellow) [GP12].

A grasp action, depicted in Figure 5.13, can be used for drag-and-drop gestures or to activate a different gesture interaction set. For example, it is feasible to perform a grasp gesture in combination with a swipe gesture to control different parts of an application. Grasp actions can be recognized depending on the object's length and width.

In contrast to pure multi-touch applications, variants employing capacitive proximity sensing have certain limitations regarding direct interaction. Considering multi-touch applications, the interaction barriers are always apparent: interaction starts when a finger touches the multi-touch surface and ends when the finger is removed. Regarding capacitive proximity sensing, the user can only make preliminary considerations about the interaction barriers, for example the height in which a hand may be detected. Thus, the interaction barriers are fuzzy and not apparent to the user.

Due to this important fact, direct feedback on the interaction status is very helpful for a user. Furthermore, it is necessary to soften Boolean decisions like selection actions. An activation state indicates when some predefined constraints are not or only partly fulfilled. These constraints can employ the vertical hand distance or a timer-controlled activation delay. Such a delay can be used for the object se-

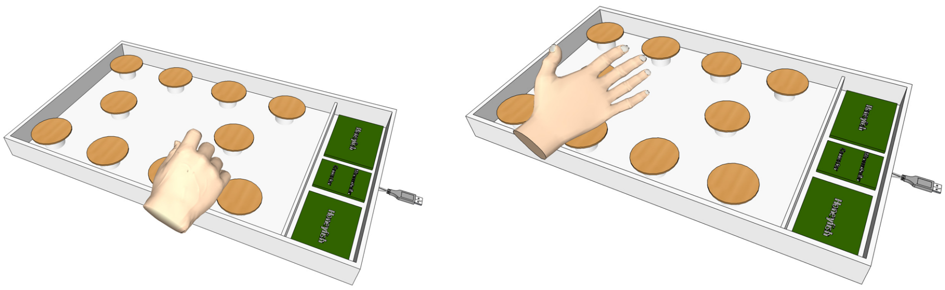

Figure 5.13.: Grasp and release actions can be utilized for drag-and-drop functionality [GP12].

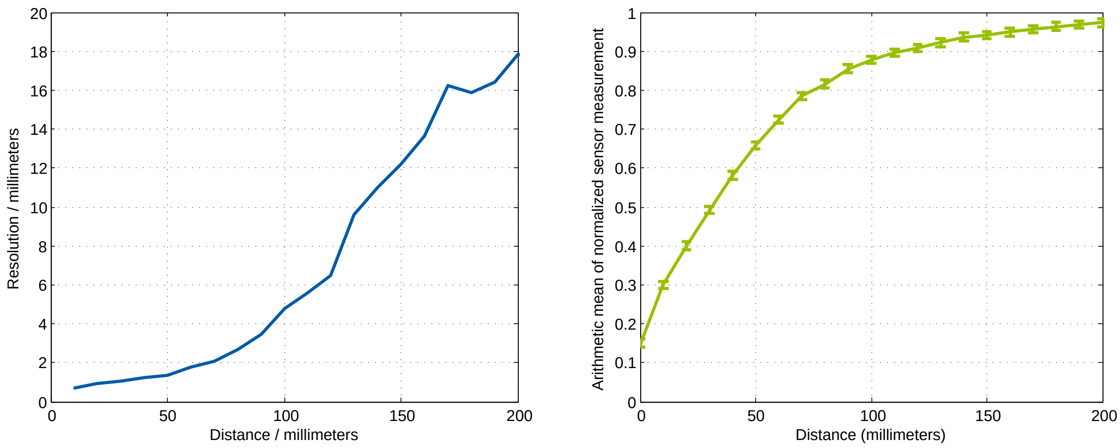

Figure 5.14.: The left plot shows a sensor's resolution in z-axis which decreases with higher object distances. The normalized sensor values and their standard variance for a constant surrogate arm distance are shown in the right plot.

lection and triggering of region-dependent actions. An exemplary activation state feedback was realized using a vertical timer bar underneath each cursor. When the cursor is not able to trigger an action, this means it is passive, it is marked in white. An active cursor, indicating that gestures can cause immediate actions, is marked in blue. Beside these gestures, region-dependent gestures were realized that trigger commands as soon as a hand remains over a predefined region. For example, this gesture type can be utilized for implementing continuous scrolling behaviors in an application. Region-dependent gestures can also be protected by a timing threshold, such that accidental movements do not cause immediate effects. Moreover, the time that a hand rests above such a region can be employed to continuously effect interaction properties, for example increase the scrolling speed with respect to the resting time.

5.2.3. Evaluation

5.2.3.1. Spatial and Temporal Performance

The object recognition performance highly depends on the hardware being used. In order to characterize the object recognition performance, we adopted a measurement setup proposed by Smith et al. [SGB99], which was later used by Wimmer et al. [WKBS07a]. The test setup uses a grounded aluminum tube acting as a surrogate arm. We took the vertical distance to the capacitive proximity sensors in relation to the acquired sensor values with their standard variance as a measure for the system's resolution.

Assuming Gaussian noise, the resolution expresses that the reconstructed distance is in a given range for 68.3% of all acquired sensor values. Figure 5.14 shows the normalized sensor values with their standard variance for the given vertical distance of the surrogate arm. In our evaluation we achieved a resolution of approximately 3.5mm at object distances around 50mm, and 35mm at object distances of 200mm. The resolution plot indicates that the resolution decreases with the vertical distance of an arm. This property of a capacitive proximity sensing system can be expressed in the object recognition model with the steepness factors γ_n.

The processing chain introduces several temporal delays caused by the hardware, particle filtering, target management, and interpolation. While the capacitive sensing hardware is able to measure the proximity to an object in realtime, the worst-case sensor update rate and the PC communication introduce a delay of approximately 30ms. With a particle filter update rate of 25Hz, new objects are usually recognized in 1-2 particle filter iterations. To suppress noisy detections at the border of the device, the target management introduces an additional delay by including new targets only after two succeeding detections. This results in a total delay of 150ms for newly appearing objects. Existing objects are tracked with delays mainly resulting from interpolation. For small movements in the area of a few millimeters, this delay is approximately 150ms in total, too, averaging over 3 succeeding object configurations. For fast movements in the area of 20cm, the total delay is only 70ms as the adaptive averaging interpolation is limited to a single object configuration, thus removing the delays associated with interpolation.

5.2.3.2. Usability evaluation

A usability evaluation was conducted at the student fair Hobit in Darmstadt (shown in Figure 5.15) with 18 participants, the majority not having a technical background. The evaluation's goal was to obtain feedback on the general user experience, including precision and reaction time, evaluate suitable applications and to compare the prototype's performance to a multi-touch system. As a prerequisite, a short introduction into the technology and the handling of the applications was given to every participant.

Figure 5.15.: Usability evaluation at the student fair Hobit in Darmstadt. The electrodes are hidden under a surface made of acrylic glass.

The first part of the evaluation focused on two gesture-controllable applications: an image viewer and a gaming application[1]. The prototype was placed on a table in front of a screen showing the application. The participants could choose to either sit or stand during the evaluation. Regarding the image viewer, each person had to accomplish a predefined set of tasks, such as image rotation, selection and browsing. The gaming application was evaluated four times - single-handed and two-handed - with one repetition

[1]Tux Racer - tuxracer.sourceforge.net

to assess the learning curve and determine if users favor either multi-hand or single-hand interaction. The collected points as well as the total time to finish a game level was recorded. In order to compare a participant's experiences with the 3D-interaction approach to a multi-touch enabled device, the same tasks had to be conducted on an ACER Iconia tablet running a standard image gallery based on Android. The reason for this comparison is the extendibility of multi-touch devices based on capacitive sensing to register 3D-interaction using the same technique. Therefore, our object recognition method can be applied as a generalized approach for interacting above a capacitive sensing device.In the second part of the evaluation, the users were asked to fill out a questionnaire to provide some qualitative feedback about their experience. The subjects had to rate their experiences on a Likert scale from 1 (no approval) to 10 (full approval). Additionally, they were asked to identify future application scenarios and advantages/disadvantages compared to multi-touch technologies.

Many test subjects experienced the evaluated prototype and its applications to be intuitive (8.71 approval) and uncomplicated. Most of them had the impression that the provided tasks could be accomplished easily (7.47 approval) and the system's reaction was comprehensive (6.94). Almost all subjects could imagine using a similar interaction device on a regular basis (8.53 approval) and deemed that the evaluated prototype is an interesting interaction modality (9.59 approval). They were fascinated by the possibility of contactless and gesture-based interaction. The large size of the prototype's interaction area was also a compelling factor. Gestures such as swiping were experienced to be recognized fast and with great precision. Both evaluated applications, the image viewer (8.0 approval) and the gaming application (8.59 approval) appealed to many subjects. Regarding the gaming application, the subjects were able to rate their favorite interaction mode on a Likert scale from 1 (single-hand) to 10 (dual-hand). Most users were either attracted to single-hand or dual-hand control, only few users liked both interaction modes. The gaming application, that was evaluated twice for each interaction mode, showed that there is a flat learning curve, letting users master the game quickly. Many users did not improve during the two rounds and achieved equally good scores from the beginning on.

The subjects had problems with the system's reaction time (5.88 approval to very fast recognition compared to very slow recognition), that can reach a maximum latency of 150ms. Especially in the gaming application, this latency turned out to be critical. Moreover, some people identified the lack of precision (5.59 approval to very high precision compared to very low precision) as an unpleasing factor. A few subjects criticized the system's recognition bounds that were not marked explicitly. Furthermore, the test persons experienced a tiring interaction posture that was caused by low table height and stretched arms. In many cases, the subjects unsuccessfully tried to influence the cursor position relatively to the current position, in a similar way as using a normal touchpad.

Compared to multi-touch interaction, the obvious advantage of contactless interaction was stated out by many people. Moreover, 3D-interaction can offer more modes of interaction. The subjects mentioned advantages like easier interaction (less fine-grained) and a seamless and invisible integration into the environment. Interaction can also be performed when wearing gloves and with less attention. On the other hand, the test persons stated that multi-touch is more precise and faster. In a direct comparison between a multi-touch image viewer and the gesture based image viewer, the multi-touch image viewer was favored by most people. This can be explained by the high interaction speed and precision that can be achieved with multi-touch technology.

The test persons were asked to identify future application scenarios. Most test persons could envision systems using capacitive proximity sensing in the area of home entertainment. In particular, the subjects suggest controlling TVs, audio, game consoles and personal computers using this technology. Due to the contactless interaction, medical and surgical applications were also mentioned very often. In contrast to multi-touch technologies, such systems might offer great advantages regarding sterility and cleanability. This is a major concern of many users who identified applications in public transport (ticket machines) and public sanitary installations. The contactless and invisible integration in furniture or behind walls and doors is an important advantage for many users. Seniors who are not able to perform fine-grained movements, for example those suffering from Parkinson's disease, can benefit from systems that recognize coarse gestures. Furthermore, conference rooms and presentation environments were proposed to be equipped with gesture recognizing systems that facilitate the interaction with such a complex technical environment.

Summing up, almost all participants liked contactless gesture-based interaction and experienced it to be intuitive and comprehensive. The precision and reaction time were criticized by several subjects. I should be possible to improve the precision by using a higher number of sensors and experiment with different volumetric representations. The latency is caused by a Java implementation running on a PC, which is currently migrated to the interaction device's microcontroller. In contrast to multi-touch technology, the interaction speed with an application is significantly slower due to longer lasting gestures. However, the strength of contactless gesture recognition lies in different application fields that cannot be covered by multi-touch technology.

5.3. Study: Object-Recognition in Front of Displays

5.3.1. Prototype

Object-recognition in front of displays has been investigated with different modalities. Based on cameras, it is possible to recognize very fine-grained finger movements in below-mm resolution [Lea14, WBRF13]. Other, less computationally expensive methods, include ultrasound [GMPT12] or infrared sensors [Mic14b]. Capacitive sensors in displays can be a low-cost alternative to those methods. The modality has already been applied in small displays on smartphones to recognize coarse 2.5-dimensional hover actions [Cyp12].

Figure 5.16.: The GestDisp prototype allows for interacting in front of a screen in distances up to 10 cm [Ber12]. On the right, an envisioned usage scenario in a car is depicted, the actual prototype is shown on the left.

In his master's thesis, Yannick Berghöfer's primary motivation was the design of a low-cost object-recognition system for automotive environments [Ber12]. Figure 5.16 depicts an exemplary usage scenario for the prototypical system, called *GestDisp*. He mainly targets the system for secondary driving tasks, which are not vital for driving security. For example, gestures can be performed for controlling music volume and the selection of songs.

In order to work towards this prospect, Yannick Berghöfer developed a prototype based on OpenCapSense which is depicted in 5.16. It was realized on an ordinary display by placing acrylic glass with transparent electrodes on top of it. The display introduces relatively high noise among many frequencies, which requires shielding in order to achieve a reasonable sensitivity. Therefore, loading mode sensing was chosen, as shielding electrodes can easily be placed underneath the sensing electrodes. But still, the noise level increases extensively when the screen is turned on. GestDisp employs eight large electrodes that cover the screen's surface. The electrodes are made of a transparent PET foil covered with indium-tin-oxide. The electrode structure was etched with hydrochloric acid onto the foils. More information about these transparent electrode materials can be found in Chapter 3.

Based on the given setup, the sensing time was adjusted to deliver 20 sensor values per second. This represents a tradeoff between spatial and temporal performance. In order to recognize a finger position, Swiss-Cheese Extended was adapted with a novel loading-mode forward model. As the eight electrodes are only able to provide a rather small set of measurements, a simple point absorber model is applied. Therefore, Monte-Carlo integration is avoided, which leads to a decreased computational effort.

5.3.2. Supported Gestures

As the system's noise is very high, a large effort on filtering the sensor values is required. Therefore, also the gestures were chosen to be very simple, mainly limited to horizontal and vertical swipe gestures. The possible gestures are depicted in Figure 5.17. Pointing gestures can be used for selecting items on the screen. They are triggered when the hand remains above the desired item for a certain amount of time. In order to raise the expressiveness of horizontal and vertical swipe gestures, they were extended by a "‘hold"’-option. In this case, the user's hand remains on the screen after carrying out the gesture.

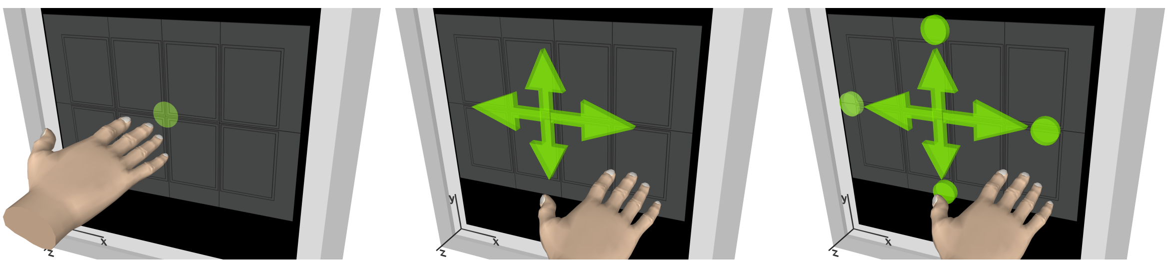

Figure 5.17.: GestDisp can be placed in front of an ordinary screen (left). It utilizes a structure of eight shielded electrodes for loading mode measurements (right) [Ber12].

Horizontal and vertical swipe gestures are used for browsing through a menu structure. Up- and Down gestures express hierarchical level changes while right and left swipe gestures mainly refer to switching items. Swipe gestures in combination with holding are applied to trigger different behaviors, such as adjusting the volume within all menu levels.

5.3.3. Evaluation

5.3.3.1. Spatial and Temporal Performance

During the prototype's design phase, a sensitive tradeoff had to be accepted between measurement time and spatial resolution. Among the reasons is the high noise level introduced by the screen which limits the gesture-recognition use-cases to simple gestures. When the screen is switched off though, a very good spatial resolution can even be obtained at smaller measurement times. In the following, I discuss how the measurement time was determined for being acceptable. 20 sensor values per second corresponds to an acquisition time of $T_{cycle} = 1/f = 50 \cdot 10^{-3}s$. Let $w = 0.25m$ be the width of the prototype's interaction space, $h = 0.15m$ its height, and $b = 0.1m$ its depth. When carrying out gestures in front of the screen, it is necessary to obtain at least $M = 5$ updated hand positions from the prototype. M depends on the characteristics of a system, with lower noise three updated hand positions can be sufficient. At hand speed $v = 0.5ms^{-1}$, the traveled hand distance between the position updates is $d_{cycle} = v \cdot T = 0.025m$. Horizontal and vertical gestures with length l can be recognized when the following constraint is fulfilled: $l > M \cdot d_{cycle}$.

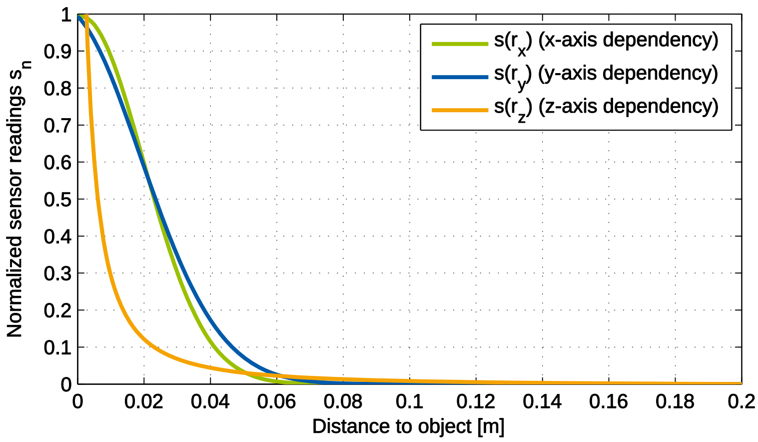

Figure 5.18.: The sensor response for the x- and y-axis is different to the z-axis. This induces the need for two types of forward model functions [Ber12].

This yields for gestures that cover the horizontal and vertical interaction space, as $h = 0.15m > 0.125m$ and $w = 0.25m > 0.125m$. However, when gestures are carried out faster or cover a smaller distance, gestures can not be recognized anymore. As movement speeds of $v = 1ms^{-1}$ are not uncommon, update rates of more than 50 Hz should be preferred. Comparing the given system to the previously discussed gesture recognition system, the interaction space is greater. Therefore, the constraints on sensor update time can be slightly lower in such larger setups.

Figure 5.18 depicts the prototype's forward reading model based on the distance of a human hand. It can be seen that the sensitivity on the z-axis (reaching out the screen) is smaller than for approaches from x- and y-axis. As described previously in the forward-reading section, two fit functions were applied to model the behavior for each z- and (x,y)-axes.

Due to the high and varying noise, the object trajectories were often disturbed at noise bursts. Therefore, additional experiments were conducted when the screen was switched off. Figure 5.19 shows three exemplary object trajectories in a plane in front of the device. The gestures lasted approximately 2 seconds, which is a rather long time compared to real-life setups. However, it can be conducted that

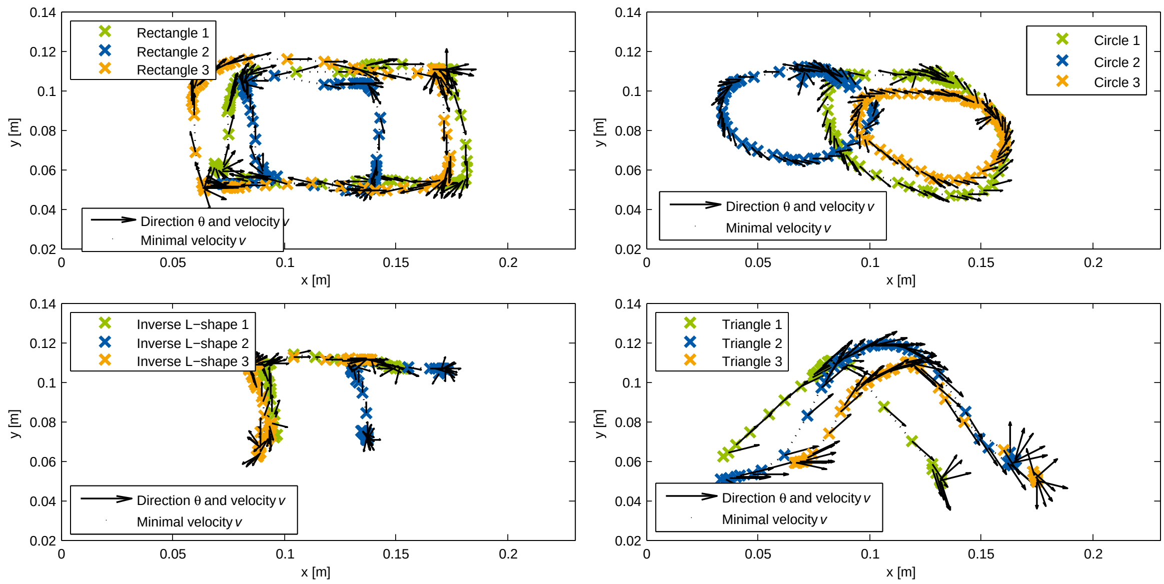

Figure 5.19.: 2D object trajectories with a deactivated screen [Ber12]. The gestures have a length of approximately 2 seconds.

the prototype enables for recognizing much more fine-grained gestures when noise can be reduced. In real-world deployments the performance can be significantly increased as all screen components can be shielded more efficiently.

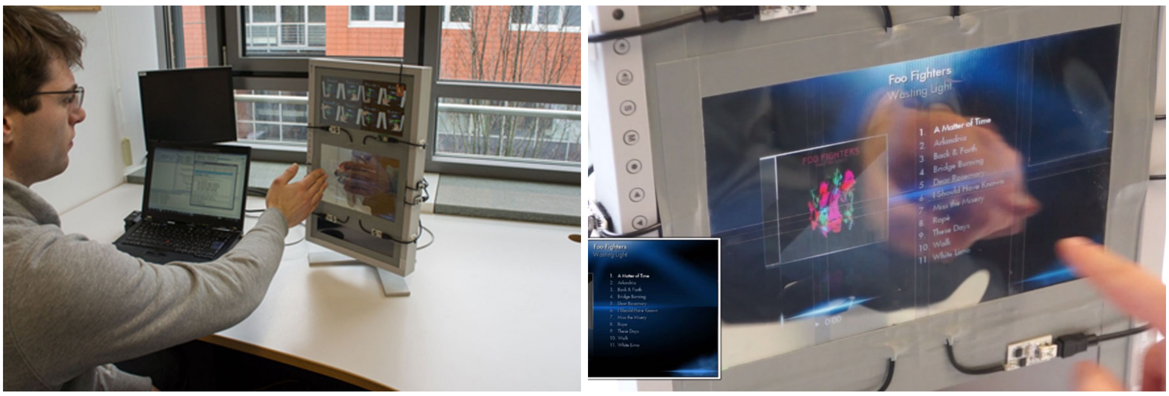

Figure 5.20.: In order to evaluate GestDisp's usability, 22 test participants evaluated a media player application [Ber12].

5.3.3.2. Usability Evaluation

Based on the prototype, Berghöfer conducted a usability evaluation to investigate the system's intuitiveness and user experience [Ber12]. 22 test persons participated in the evaluation with an average age of 32.31 years. These persons evaluated an infotainment application intended to be used within automotive environments. Most participants were used to gestural interaction paradigms and employ them on a regular basis (77%). The majority of 18 participants drove a car regularly. The participants had to accomplish a set of tasks by navigating through a menu structure. This included selecting songs in a music library, and switching radio stations. Moreover, the participants were given a short cheat-sheet on the performable gestures. In the following, questions based on a on a Likert scale from 1 to 9 were asked to determine the prototype's usability.

The learning curve for the gestures was very steep after getting used to the system. But still, people were given instructions on how to perform certain gestures. The problem of learnability and intuitiveness of gestures is discussed later on in Chapter 7. Although the performable gestures were learned by a cheat-sheet, the linkage between gestures and actions was perceived as reasonable (7.86 approval). Most persons agreed that the tasks were easy to execute (7.50 approval). Very similar to the previous evaluation, some participants criticized the reaction time and reliability of the system. The reliability was mainly influenced by noise, which caused false object detections leading to a system reaction. Most of the participants stated that the system's reaction was transparent (7.77 approval). The evaluation also discussed the general applicability of the presented use-case in automotive environments. Many drivers were less satisfied with established media player applications in cars (5.86 approval to totally satisfied). Interestingly, the approval turned out to be even less for co-drivers (5.50 approval).

Berghöfer also asked open questions referring on the general usability. Similar to the previous study, four people mentioned that in-the-air gesture recognition is useful in automotive scenarios. This goes hand-in-hand with the lack of haptic feedback, which was mentioned as a negative property by four persons. During the study, it became obvious that many persons suffered of a tiring arm after interacting for a while. The intuitiveness was evaluated positively by nine participants. Among the reasons are the strong correspondences to multi-touch interaction. One participant mentioned that this could be a problem for the elderly, who might not yet be familiar with multi-touch technology. Four participants mentioned that the learning effort is very high, supporting the fact that a system should give clues on possible gestures and their outcomes.

5.4. Summary

In this chapter, I presented *Swiss-Cheese Extended*, a method for object recognition based on an array of capacitive proximity sensors. The method allows to recognize multiple object parameters and can be easily adapted to different kinds of volumetric objects. Based on particle filtering, objects can be recognized and tracked in real-time. Using a time-series of object locations, it is possible to perform gesture recognition, for example with heuristics, HMMs, or DTW. The method was evaluated with two exemplary systems, relying on shunt mode and loading mode sensing.

The first study describes a custom-built gesture-recognition system. Here, Swiss-Cheese Extended is applied to track the state of two hands using just 16 sensor channels. The system supports various gestures, including multi-hand rotation and zoom as well as grasping actions. These gestures allow controlling different demonstration applications that were adapted for the input device. To evaluate the performance and user experience, a study compared our system to a multi-touch tablet. While the test persons considered the system improvable in terms of interaction latency and precision, most participants valued the intuitiveness and novelty of the device.

The method was also evaluated on a prototype for interacting in front of a display [Ber12]. Due to the limited resolution of only eight sensors, a unit absorber model was applied for modeling the finger position. The sensing mode's different behavior induced the need for incorporating a new model for loading mode measurements into Swiss-Cheese Extended. HMMs were applied for recognizing gestures, which were mainly limited to easy swipes. A usability evaluation was performed that supports the intuitiveness of the system. Very similar to the first study, interaction speed and reliability was criticized.

The primary reason is the noise induced by the underlying display and mechanical transformations due to varying temperature.

Among both studies, various potential application scenarios were identified, ranging from medical solutions, where sterility is crucial, to unobtrusive integration in furniture. Particularly in assistive applications, the presented method can enrich the execution context, for example by identifying postures in beds. This requires incorporating a more complex volumetric model that models the human skeleton. Gesture recognition in front of displays is applicable in public transport, show windows, and museums. The method can be applied in a very generic way on different sensing problems in the capacitive domain. Further investigations on forward models and volumetric object representations can lead to reasonable extensions of the method. Moreover, it would be highly interesting applying the method to different ranging modalities, such as ultrasound, Lidar, and radar.

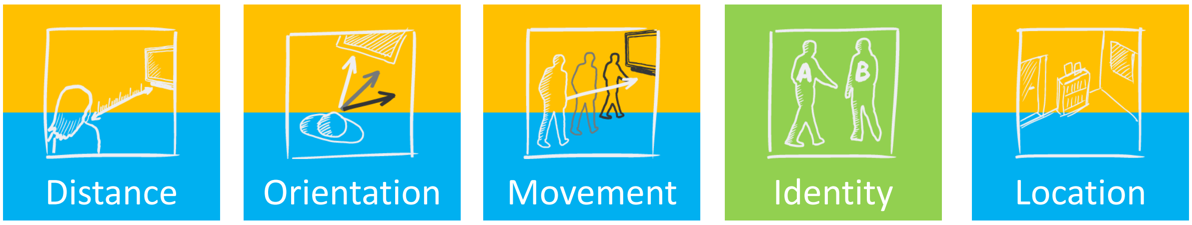

Figure 5.21.: Swiss-Cheese Extended completes the picture of proxemic interaction dimensions in this thesis [GMB*11].

With Swiss-Cheese Extended, I was able to tackle the second research goal of interpretation and fusing of data obtained by capacitive proximity sensors. Figure 5.21 depicts the proxemic interaction dimensions that are covered by Swiss-Cheese Extended. Combining my contributions in physical sensing opportunities, namely OpenCapSense and CapNFC, with Swiss-Cheese Extended enables me to complete the picture. Based on the investigated proxemic interaction modalities, it is now possible to allow for a detailed environmental perception. Herewith I conclude the discussion on the field of *Environmental Perception* and move to *Ubiquitous Interaction*.

6. Context-Aware Devices and Environments

In the previous chapters I discussed three aproaches to capture the dimensions of proxemic interactions for *Environmental Perception*. Based on the retrieved physical data and its interpretation, I am able to retrieve measures on distance, orientation, movement, identity, and location. As shown in Figure 6.1, I will now investigate the exploitation of these dimensions for *Ubiquitous Interaction*. Here, I differentiate between implict and explict interactions. In the domain of explict interaction, a user interacts with a device by intentionally giving more or less abstract commands. Explicit interaction comprises settings like pressing buttons on a touchscreen or carrying out gestures above a sensing surface. In implicit interaction though, a system understands a user's actions even if they are not intended for direct interaction [Sch00].

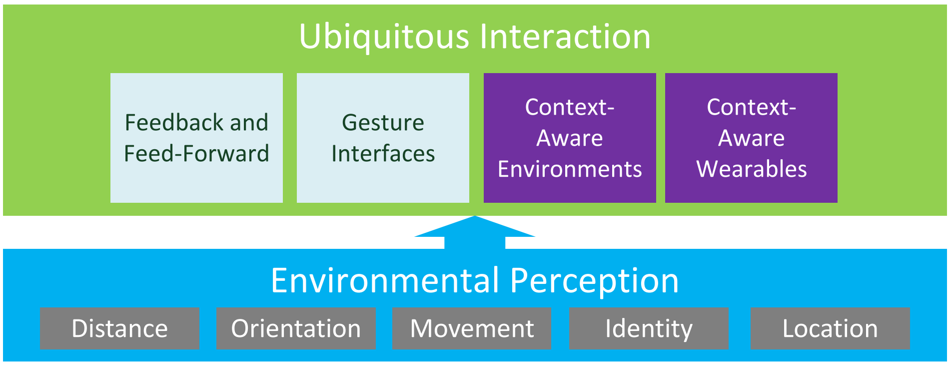

Figure 6.1.: The focus of this chapter lies on the implicit part of Ubiquitous Interaction. It includes wearable as well as stationary appliances.

Implicit interaction systems can sense the user's situation in sufficient detail to respond intelligently to the current situation. In [Sch00], Schmidt underlines the necessity of perceptional capabilities to understand the user's situational context: "We will be able to create [..] devices that can see, hear and feel. Based on their perception, these devices will be able to act and react according to the situational context in which they are used.". Context or context knowledge can be regarded as an abstract representation of the user's situation within a computing system [Sch00]. The concept of implicit interaction matches very well with the vision of UbiComp [Wei99], as it can be seen as one of the enabling factors. A detailed understanding of the user's context is required to intelligently sense a user's goals and give intelligent support to achieve them. Moreover, the disappearance of computing abilities requires implicit interaction as systems inherently become visible during explicit interaction [SVL01].

Recognizing a user's activities can contribute to a detailed understanding of a user's situational context. One aspect in the wide field of activity recognition is the recognition of physical activities [BI04]. In this chapter, I will show how capacitive sensors can contribute to context knowledge by recognizing physical activities. The unobtrusive placement of capacitive sensors makes them ideal to be embedded in everyday objects and furniture. Besides those often static objects, capacitive sensors in wearable devices can also reveal meaningful information about the environment. Figure 6.2 de-

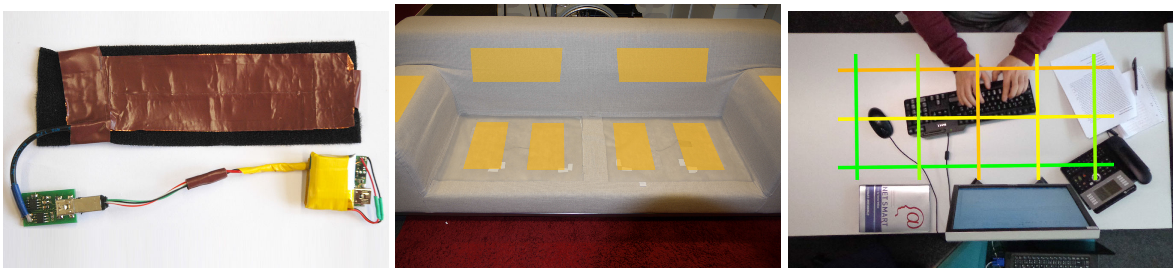

Figure 6.2.: Capacitive sensing to derive the situational context of a user based on recognizing physical activities. I present three use cases from a wearable device (left), to a posture-recognizing couch (middle) and tabletop (right).

picts the three different use-cases that are presented in this chapter. This chapter is based on three papers [GPMB11, GPBB12, GPBK*13], which will be introduced in the beginning of each section.

6.1. Wearable Capacitive Sensing

This section is based on the paper [GPBB12]. The use of 'we' corresponds to the authors Tobias Grosse-Puppendahl, Eugen Berlin, and Marko Borazio.

Sensing a person's activity is an active research topic with a raising interest due to the advancement in mobile phone technology. These devices include multiple sensors and therefore enable the recognition of daily activities [BGC09]. Current wearable activity recognition systems are able to unobtrusively capture and recognize a person's activities throughout the whole day. These systems often rely on inertial sensor data that is captured by wearable sensors embedded in a mobile device [BGC09] or attached to the body [RDML05]. Usually, single sensor modalities are used or duplicated to detect the activities. However, it is a great challenge to identify fine-grained activities just by using a single modality like the accelerometer.

Activity recognition research relying on wearable sensors mostly considers inertial data from the participants body to infer performed activities, such as in the works of [RDML05, BI04, SCC10, ASLT05]. The acceleration data is often augmented with data from sensors such as gyroscopes [HSAT10], magnetometers [AB10], ambient light [BV12] or ambient and skin temperature [KSSF03], aiming to extract a more detailed environmental user context. In [WM10], the authors use heart rate information as an addition to the accelerometer data to detect activities like lifting and lowering loads or even digging. In [WLTS06], workshop assembly activities are detected by augmenting acceleration sensors with microphones.

The works of Fishkin et al. [FPR05] and Patterson et al. [PFKP05] show that detecting touched and used objects can be very helpful for activity recognition. By using RFID readers embedded in gloves or bracelets at the wrist and RFID tags attached to various objects of interest, one can detect the object grasped and used by the user. This enhances the activity recognition in various application scenarios, such as activities of daily living [SHVLS08b] and [PFP*04], activity tracking in car manufacturing [SRO*08], or household and gardening activities [BLvLS10]. Cheng et al. have investigated the possibility of using capacitive sensors for activity recognition by measuring shape changes of muscles and skin [CAL10]. To our knowledge, capacitive proximity sensors have not been embedded into a wearable device to enhance the performance of activity recognition.

Capacitive proximity sensors can indirectly measure the distance and nature of a grounded object within reach. This means that the measurement result depends on the object's distance, its size and the material it is made of. In this section, we show that an accelerometer and a capacitive proximity sensor can be used to improve activity recognition in activities of daily living. Therefore, we obtained an open-hardware and open-source wrist-worn activity data logger [hed] and integrated a capacitive proximity sensor invisibly into the wristband.

There are several intuitive examples for which a combination of accelerometer-based activity recognition with a capacitive proximity sensor reveals its strength. For example, it may be conducted which material is placed underneath the wristband. The capacitive proximity sensor would return a different measurement result for a hand placed on a couch covered with fabric than a hand placed a wooden table. Moreover, the approximate distance from the wristband to objects can be exploited to identify activities like grasping into a locker or a refrigerator to prepare food.

Our approach to enhance the acceleration data from a wearable sensor is comparable to the RFID scenarios just mentioned. It also relies on a single wearable sensor and an unobtrusive deployment. The main difference lies in the fact that we do not consider an accurate detection of tagged objects, but the proximity and nature of unknown objects in the environment. We apply loading-mode capacitive sensing where a single electrode builds up an electric field to any grounded object in the environment. By measuring the capacitance, conclusions can be made upon the proximity and nature of an object. Using loading-mode is also advantageous considering the complexity of the system. It requires only a single shielded electrode that can be integrated invisibly into the wristband.

Especially in wearable scenarios, capacitive proximity sensing faces the great advantage of being robust against changing lighting conditions and occlusion. Moreover, sensing electrodes can be integrated invisibly into wristbands or clothing. On the other hand, the exact distance to objects can only be approximated since the object's surface, its conductivity and grounding has influence on the measurement result. A single sensor will thus deliver data that has a certain degree of ambiguity. Due to the nature of capacitive proximity sensors, they can be prone to errors in environments with strong and rapidly changing electric fields. This, however, is usually not an issue when considering activities of daily living.

6.1.1. Hardware

This section presents the two components of our hardware prototype, the wrist-worn activity data logger tailored to capture acceleration data, and the capacitive proximity sensor used for distance measurements. The HedgeHog sensor [hed] is a custom designed wearable data logger aiming at long-term deployments in activity recognition scenarios. Due to its small form-factor (37x32x16mm) and weight, this wrist-worn sensor is an unobtrusive way to record relevant motion data.

The sensor node itself is built around the low-power Microchip microcontroller (PIC18F46J50) featuring an accelerometer sensor (ADXL345) to capture human motion, light and ambient temperature sensors and a microSD flash card for locally storing the sensor data. The sensor is powered by a 200mAh lithium polymer battery, which allows for two weeks of continuous recording on a single battery charge. A USB port is used to configure the sensor (e.g. setting the sensitivity of the accelerometer), to access the stored sensor data, and to recharge the battery. A plastic case nicely packages and protects the sensor to be worn at the wrist (Figure 6.3).

The 3D accelerometer sensor is being sampled at 100Hz, resulting in 10ms equidistant measurements. For efficiency reasons, the sensor data is run-length encoded before being stored locally to the microSD card. The HedgeHog can be extended with further sensors tailoring different application scenarios. For our scenario, we have added a capacitive proximity sensor that is described in detail in the next section.

Figure 6.3.: The inertial data logger featuring a low-power microcontroller, a 3 axis accelerometer, a microSD flash card for storing the sensor data and a USB connector for accessing the data (on the right) is powered by a small lithium polymer battery and is packaged into a plastic case to be worn at the wrist (a version with an OLED display).

A wrist-worn capacitive proximity sensor requires a shield that eliminates the influence of the grounded arm directly underneath the sensor. Using this setup, we can detect the proximity to a grounded object in the environments for distances up to 20cm. Especially for mobile devices, it is required that the sensor draws a very small amount of power. Thus, other proximity sensing input modalities like ultrasound or optical measurements are not applicable for this type of mobile application.

The capacitive proximity sensor performs measurements in loading mode. Two electrodes are integrated into the wristband, one sensing electrode and one shielding electrode. The sensor draws a supply current of 1mA at 3.3V when active and qualifies it for wearable proximity sensing applications. In the following a virtual capacitor denotes the capacitance between the sensing electrode and the environment. The capacitance of the sensing electrode to environmental objects increases with closer distances.

The sensing circuit is taken from OpenCapSense, presented in Chapter 3. It is based on a timer that controls the charging and discharging cycles of the virtual capacitor that is built by the sensing electrode and the surrounding environment. The timer toggles from charge to discharge at the time when a threshold voltage at the capacitor is reached. This results in an astable operation with succeeding charge/discharge cycles. When the capacitance of the virtual capacitor increases, the charging time will also increase and vice-versa. Therefore, the capacitance is inversely proportional to the number of charging cycles in a given time span. In order to guard the sensor from measuring the capacitance to underlying objects, a shield electrode is placed directly underneath the measuring electrode. The shield is driven with the same potential as the sensing electrode, such that the capacitance between the two electrodes is negligible. Using this shielding method, the measured capacitance will only be slightly affected by the grounded underlying arm. Figure 6.4 shows the final wrist-worn prototype used in the evaluation experiments, with the HedgeHog as the main data logger. The capacitive sensor circuit and the wristband holding the sensing and shielding electrodes.

The operations required for a measurement cycle are illustrated in Figure 6.5. The proximity sensor board generates a clock signal with varying frequency depending on the charge and discharge cycles. The HedgeHog measures the resulting capacitance by counting the signal's edges over a gate time of approximately 9.5ms. During that counting phase, the microcontroller is sent to sleep in order to reduce

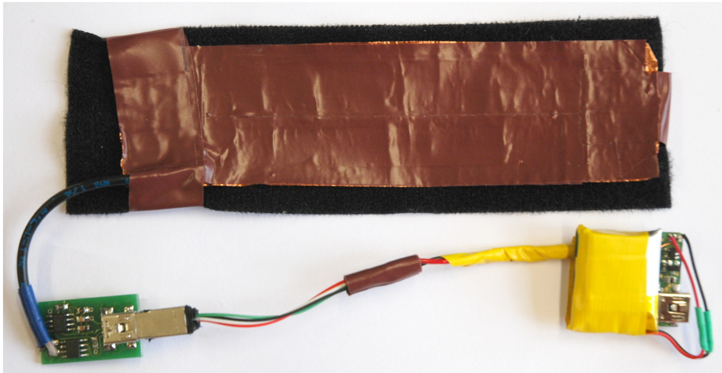

Figure 6.4.: The hardware prototype at a glance: HedgeHog activity logger at the lower right, the capacitive sensor unit at the lower left, and the wristband with the sensing and the shield electrodes on-top each other. The electrodes are covered with adhesive tape for isolation purposes.

power consumption. In the following, the HedgeHog applies run-length encoding on the measured data to reduce overhead and periodically logs the data to the integrated microSD card.

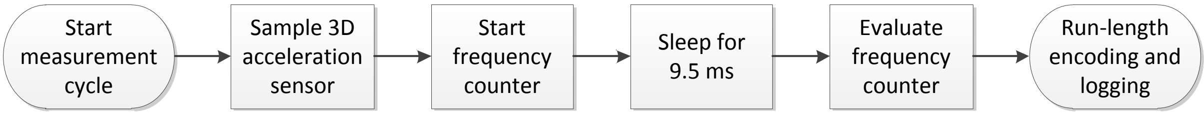

Figure 6.5.: Overview of the measurement procedure carried out by the HedgeHog sensor: using the microcontroller's Timer0 module in counting mode, the oscillating signal generated by the capacitive sensor circuit can be measured by counting the frequency pulses over a predefined gate time of approximately 9.5ms.

6.1.2. Experiment

This section presents the experimental setup including the activities and the participants, as well as the findings that were obtained during the evaluation. The experiment setup aims to depict a typical scenario of a person in daily life. Especially in the field of Ambient Assisted Living (AAL), it is desired to monitor activities like drinking, preparing lunch and sleeping. A fine-grained monitoring of such activities may help elderly or people suffering from mental diseases to maintain a healthy day/night rhythm and take action if irregularities occur. Figure 6.4 shows the modified HedgeHog activity logger that has been extended with a capacitive proximity sensor. The wristband has two electrodes, a sensing electrode underneath a slightly bigger shield electrode.

The recorded test set contains the following activities: opening door, sitting on a couch, lying on a couch, putting kitchen equipments from a shelf and out of a locker, preparing a bread with marmalade, eating the bread, pouring and drinking water, walking and sleeping. The relations of these activities to environmental objects are given in Table 6.6. Some of those activities are very hard to recognize when the data is limited to a single modality like a 3D accelerometer. For example, sitting at the table and sitting on a couch are very similar activities. We aim to show that the data basis can be significantly improved by the additional input modality.

In order to evaluate if capacitive proximity sensors in wrist-bands can enhance the performance of activity recognition, we have conducted an evaluation with 7 test persons. All test persons received

activities	objects involved	objects nearby
open door	door knob	door
sitting	chair or couch	body, chair, couch, table
lying	couch	body, couch, cushion
get things	plate, glass, cutlery, bread, marmalade, bottle	shelf, locker, fridge, table
making bread	bread, knife, marmalade	table, plate
eating	bread	table, plate, body
drinking	bottle, glass	table, body
sleeping	bed, cushion, blanket	body
walking		body

Figure 6.6.: Some details on the activities performed during the experiment and objects directly involved or nearby.

a basic plot with the activities they were supposed to perform. They were not given any instructions about the way they are supposed to perform the activities. After manual labeling, we used this test-set as ground truth and performed a 4-fold cross-validation on an support-vector machine (SVM) classifier on each user. The cross-validation was performed once with and once without including the data of the capacitive proximity sensor into the feature set. We chose an SVM classifier because of its high relevance in activity recognition and its fast performance.

The classifier was trained with basic features that were extracted from a sliding window of 1 second width. Our first tests have shown that greater window sizes do not provide better classification results. In order to suppress noise contained in the capacitive proximity sensing data, we applied a moving average filter with a kernel size of 10. The final feature set contained the arithmetic mean, min, max, median and standard variance for each accelerometer axis and the capacitive proximity signal. These simple feature types represent standard features applied in activity recognition. Since we aim to show an improvement using the new modality, the selection of features and classifiers does not represent the primary focus of this paper.

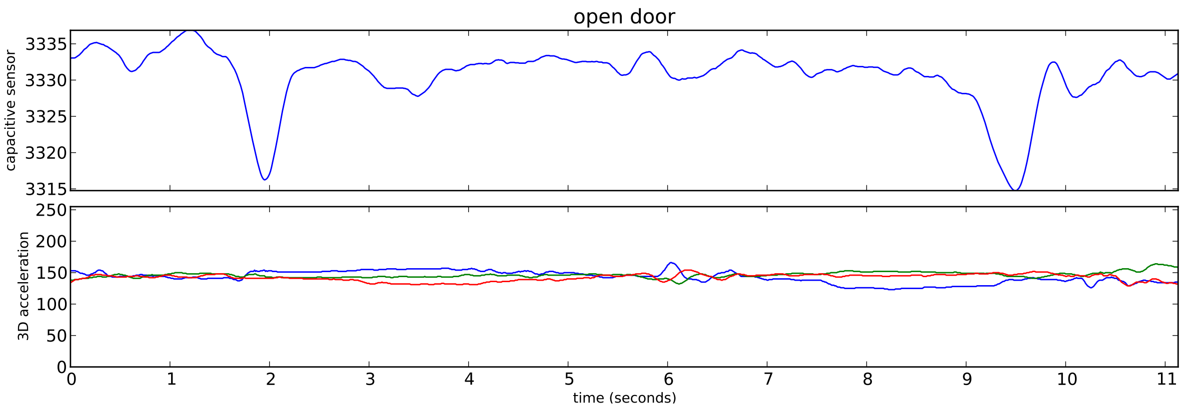

Figure 6.7.: When the participants entered the apartment, the wristband approached the door knob twice, at the time of opening and closing the door. This fact can be observed in the capacitive proximity data (upper plot) at the beginning and at the end of the activity, whereas the acceleration has no characteristic information (bottom plot).

In the following, the performed activities will be analyzed in detail, stating out the influence of the capacitive proximity sensor on the classification result. In general, the usage of data provided by the new modality showed improvements in recognition rates reaching from 2.4 up to 10.7% for a single activity.

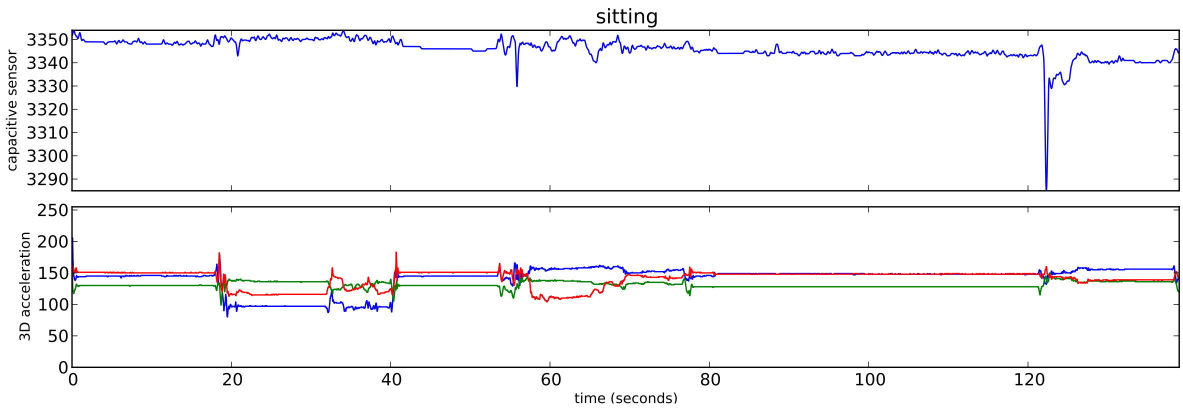

Figure 6.8.: Example of the "sitting" activity in which the user moved his hands quite frequently (bottom plot). Most of the time the values of the proximity sensor stay more or less constant, probably due to the hands position on the couch's fabric. The sharp peak in the capacitive sensor data (upper plot) occurred when the participant scratched the back of his head.

The opening door activity has very poor recognition rates without the data from the proximity sensor. The average F-measure could be increased from 35.5% to 46.2%. A plot of the activity is given in Figure 6.7. The capacitive proximity sensor shows two approaches to the door knob, one for opening the door (2s) and one for closing the door (9s). The acceleration sensor captures relevant data in the time in which the person moves into the room and the hand changes from the outer to the inner door knob (5 - 7s). The recorded data for this activity also shows strong correlations between all experiment participants. The confusion matrices show that the "open door" activity was often confused with the "sitting" activity, probably because of the amount of motion on the one hand and the proximity to nearby objects (door, couch or cushions) on the other hand. By using the capacitive sensor data, the recall for that class and confusion with the sitting activity could be improved.

After closing the door, the participants were supposed to sit down on the couch. It turned out that there are great variations of the sitting posture and the corresponding hand positions. Many users tapped with their fingers or hands while sitting, changed their sitting positions very frequently, or were even talking and gesticulating, as shown in Figure 6.8. In this case, it is obvious that the data from the acceleration sensor is very difficult to interpret as there are numerous changes in the axial orientation of the sensor. However, the capacitive proximity sensor is able to indicate when a hand is placed on the surface of the couch. Especially for this particular participant, the F-measure increased from 50.3% to 60.4%, while the average F-measure improved from 68.0% to 74.4%.

In the following, the participant were instructed to lie down on the couch. Again, there were great variations in how this activity was performed by the participants. For example, some of them crossed their hands under their head, or placed them on their body. For this class, the average F-measure could only be increased by 2.4%, from 81.2 to 83.6%. Considering some participants, the activity was often confused with the "making bread" class. By using the proximity modality, the confusion between the two classes could be reduced.

After that, the participants were asked to walk over to the kitchen and to put food and dishes from a shelf and a locker on the table. This activity involved direct interactions with various objects as well as proximity to furniture in the room (see Table 6.6). The capacitive proximity sensor was able to capture the proximity to the shelf and to the table. The average F-measure for this activity is rather low, but improved by 6.8% from 53.8% to 60.6%. The worst performing participants for this activity

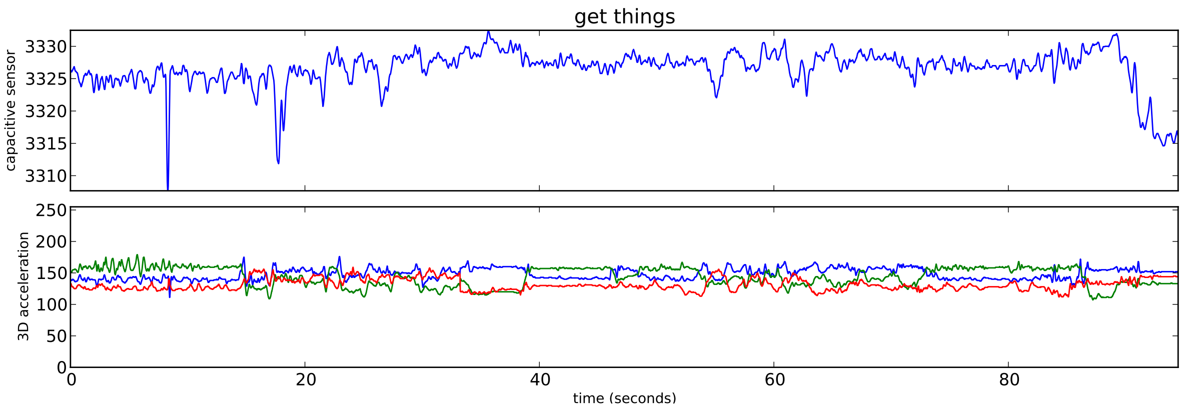

Figure 6.9.: An example of the "getting things" activity, where the participants had to get food and dishes from shelves and lockers. The proximity sensor peaks in the beginning (9s and 19s) indicate immediate proximity to shelf, and to the locker (55-63s) in the kitchen (upper plot). The signal drop at the end results from the participant placing his hand on the table when she was finished.

reached an F-measure of 46.8% without and 53.5% with the capacitive sensor, while the best performing one reached 62.0% and 64.5% respectively. The low performance results from confusions with other activities, with a higher tendency to the "prepare bread" activity across all participants. This is most likely due to various objects involved in both activities. The capacitive sensor modality has a much more positive impact reducing the confusion with other activities.

Figure 6.10 shows an example instance of preparing a bread with marmalade. It is noteworthy that the acceleration data does not seem to provide any characteristic patterns, while the proximity sensor indicates a table, plate, or other objects in immediate distance. This activity showed a high improvement in the average F-measure by 10%, from 49.0% to 59.8%, where the data delivered by the capacitive proximity sensor is taken into account. The making bread class was often confused with the sitting class for some users, probably due to lots of motion during the sitting, as mentioned previously. For other users, making bread was confused with eating or drinking. Using the new input modality, confusion across users could be reduced.

When considering the eating activity, the impact of the new capacitive proximity sensor on the classification performance is quite low. Some of the participants ate their bread leaving their hand close to the mouth, while others moved their hand up and down putting their bread aside on the plate. This results in the performance range from 71.2 to 88.1% without and 77.5 to 90.6% with the proximity data. An example of an eating activity is shown in Figure 6.11, where the participant took a few bites from the bread while putting it down every time. The average F-measure for that activity increased slightly by 2.7% (from 79.5% to 82.2%).

The activity "drinking water" is depicted in Figure 6.12. The participant took a few drinks from the glass, while leaving the hands lying on the table in between. These motions can be easily detected in the acceleration as well as proximity data. The accelerometer data shows that there periodic up- and down-movements while the capacitive proximity sensor delivers data that is associated to the proximity of the table. The average classification performance lies at 48.4% without and 54.4% with the proximity data taken into account, resulting in a gain of 6%.

Regarding the walking activity one can identify periodic changes in the acceleration as well as in the measured capacitance, illustrated in Figure 6.13. While walking, the capacitance between the wristband

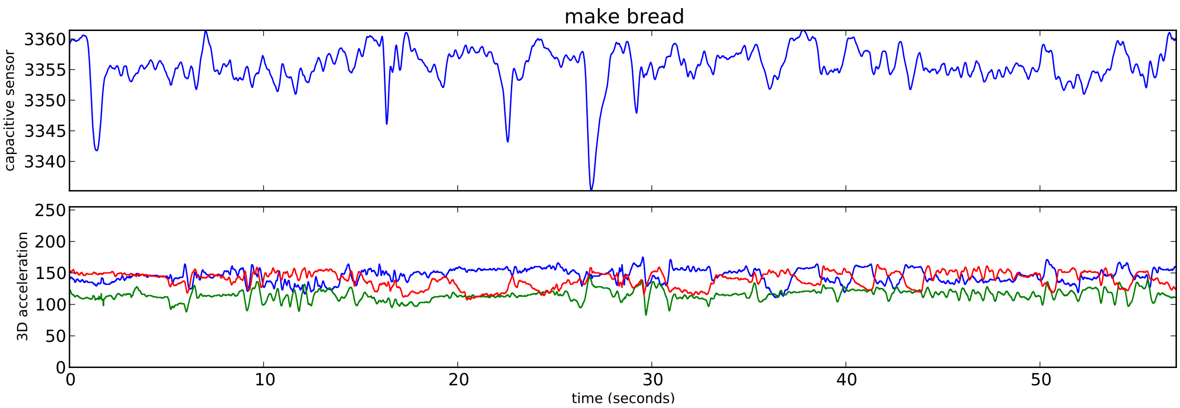

Figure 6.10.: An example of the "preparing bread" activity, where the participants had to put marmalade on a slice of bread. The proximity sensor indicates the closeness to the table, while the acceleration sensor shows recurring hand motions.

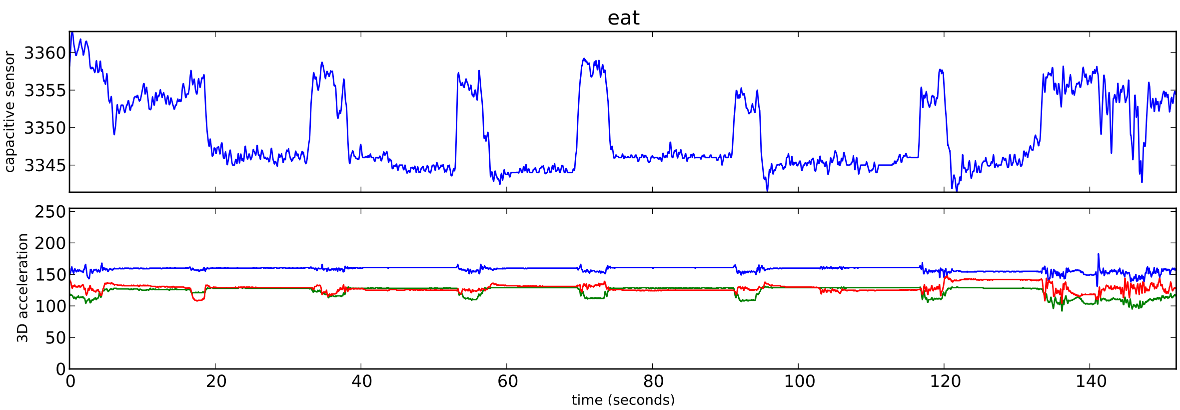

Figure 6.11.: An example of a participant eating a slice of bread with marmalade, taking 5 bites from it. After each bite, the hand is placed on the table, which can be recognized both in the acceleration as well as the proximity data plots.

and the leg increases when the wristband is located close to the body and decrease when the wristband moves away. There were problems distinguishing this activity from "get things" that could be improved by using the new input modality. The classification improvement for this activity accounts to 12.2% boosting the average F-measure from 41.7% without to 53.9% with the new sensor. Due to the low performance results and the characteristic periodic signal shape, it would help to consider frequency domain features, as it is often applied in related work.

The sleeping activity (an example shown in Figure 6.14) was classified with an average F-measure of 83.2%, which increased to 86.9% when using the proximity data. In this case, the capacitive proximity sensor is able to capture the surrounding cushions and blankets, as well as the body or head of the participant. The accelerometer data and the proximity data has larger periods with a constant signal, which is a cause for confusion with the "lying on the couch" activity.

Figure 6.15 depicts two confusion matrices for an exemplary participant from our evaluation, once without and once with including the proximity data. In most activities, an enhancement in the number of correctly classified instances is observable. The "lying" activity's recognition performance could benefit a lot from the proximity data, improving both precision and recall. A better classification performance

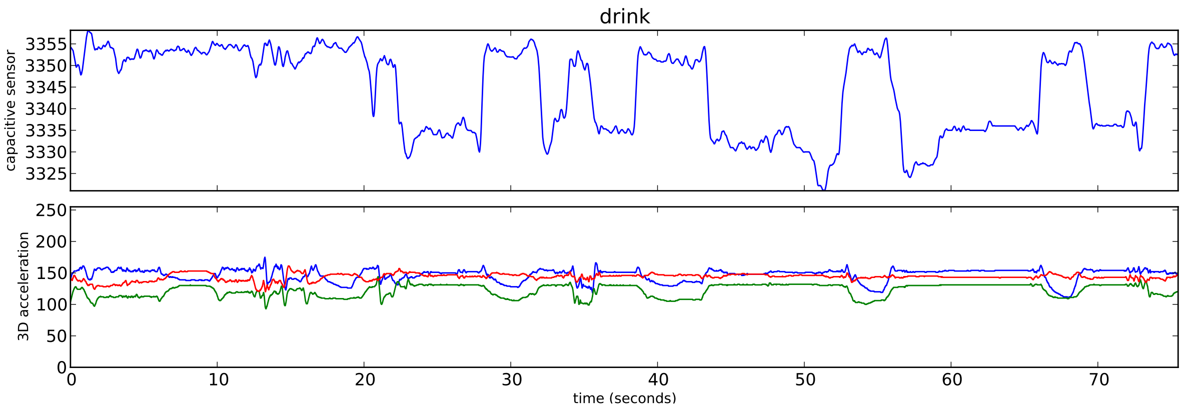

Figure 6.12.: An example of the “drinking” activity. The participant first pours some water into the glass and then takes three drinks of water. After each sip, he returns his arm to the table which can be observed in the characteristic patterns of the proximity sensor.

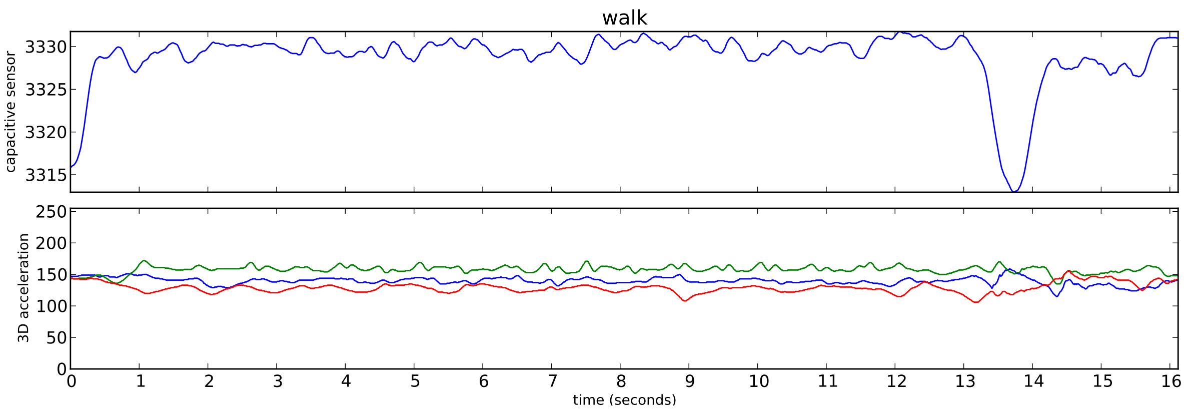

Figure 6.13.: An exemplary instance of the class “walking”. The acceleration sensor and the proximity sensor show periodic recurring patterns that are related to the pendulum-like arm movement and the proximity to the person’s body during those movements.

can also be observed for the “get things” class that includes interactions with a multitude of objects in the environment. Regarding the activity “make bread”, the capacitive proximity could reduce the number of confusions with the “eating” class significantly.

Due to the high similarity of eating and drinking (cf. Figure 6.11 and 6.12), the number of confusions between those two classes increases when considering the proximity data. However, for the “drinking” class the new modality limits the confusions to related activities such as “eating” and “make bread” only, while lowering the number of false recognitions for the other activities.

In this section, we presented a wrist-worn activity data logger prototype, which consists of an accelerometer in combination with a capacitive proximity sensor integrated into the wristband. Our experiments with seven participants and nine basic daily activities show that this additional input modality can significantly boost the activity recognition performance. Regarding the classification performance, we obtained an improvement in the average F-measure of 6.3%, from 67.2 to 73.5%. Specifically, the activity classes “walking”, “preparing bread” and “open door” could benefit a lot from proximity-related sensor data. For such classes, the classification performance could be boosted by 12.2, 9.0 and 10.8% respectively.

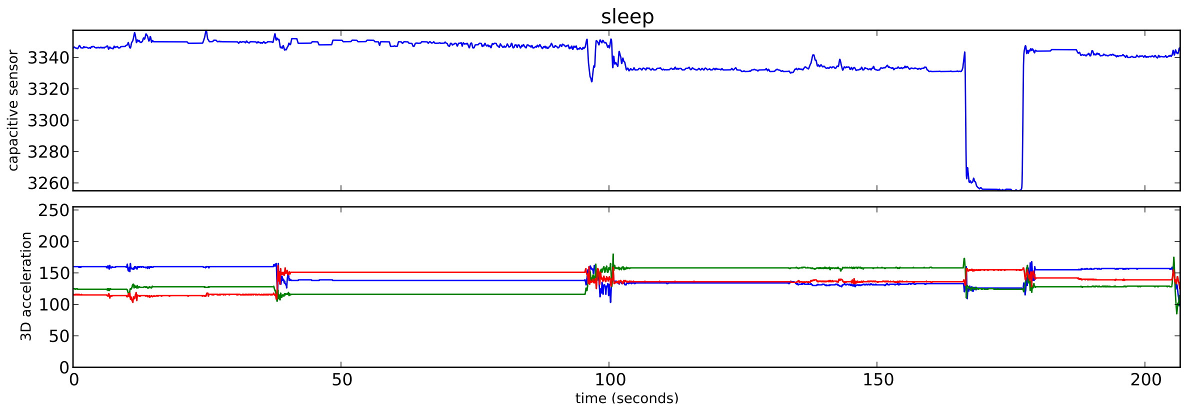

Figure 6.14.: During the sleeping activity the data from both sensors remains constant for large time spans. One can draw conclusions about the coverage of the arm with either cushions, blankets or the proximity to the mattress, the head or body of the participant.

without proximity data									with proximity data									
a	b	c	d	e	f	g	h	i	a	b	c	d	e	f	g	h	i	← classified as
8	0	0	0	0	0	0	2	1	**11**	0	0	2	1	0	0	0	1	a = open door
0	**140**	0	5	2	0	0	0	0	0	**140**	0	5	2	0	0	0	0	b = sitting
0	0	**294**	1	9	10	1	1	0	0	0	**309**	1	0	4	0	0	1	c = lying
1	7	1	**50**	17	9	1	4	1	2	6	0	**67**	14	9	1	2	1	d = get things
0	2	7	7	**102**	42	5	0	2	1	1	0	13	**120**	24	7	1	1	e = make bread
0	0	3	5	21	**288**	4	0	0	0	0	2	1	19	**300**	3	0	0	f = eating
0	5	11	1	29	40	**18**	0	0	0	4	4	0	26	57	**16**	0	0	g = drinking
2	0	0	2	0	0	0	**11**	0	2	0	0	0	0	0	0	**14**	0	h = walking
0	2	2	0	2	0	0	0	**271**	0	0	0	0	1	0	0	0	**276**	i = sleeping

Figure 6.15.: Activity recognition evaluation revealing the positive impact of the capacitive proximity sensor. Here, we are comparing SVM classification presented as confusion matrices for an exemplary user, without the proximity data on the left, and with the proximity data on the right. Note that the reject class (background data not annotated as an activity) is not included in the confusion matrix.

With this proof of concept we show that the proximity information can provide an information gain regarding the evaluated activities. In future work other relevant feature types should be investigated, such as frequency domain features, as well as feature sets to extract the most discriminative ones. Additionally, using other classifiers (such as HMMs) or more fitting classifier configurations might also improve activity recognition performance.

The classification results could also be improved by using more than one sensing electrode in the wristband. For example, the wristband could integrate up to four electrodes that are placed on each side of the arm. However, this will lead to smaller electrode surfaces thus resulting in a decreased sensing distance. The sensor's power consumption can be decreased by shorter measurement windows and the choice of more energy efficient hardware components as well as software implementation. Measuring pulse width lengths instead of counting the number of pulses of the sensor's signal may reduce the required time needed for a measurement, thus increasing the time the microcontroller is able to sleep.

Capacitive proximity sensors represent a suitable new input modality for future activity recognition systems. The low power consumption as well as the invisible integration of a sensor into the wristband meets an essential requirement of wearable and unobtrusive applications. Especially in AAL environments, these systems can help monitoring the course of chronic diseases by recognizing activities of daily life. This may improve the quality of life of persons affected and their caregivers.

6.2. Posture-Recognizing Furniture

This section is based on the two papers [GPBK*13, GPMB11]. The use of 'we' corresponds to the authors Tobias Grosse-Puppendahl, Andreas Braun, Alexander Marinc, Sebastian Benchea, Felix Kamieth, and Christian Schuster.

As stated in the introduction of this chapter, understanding a person's situational context is essential for intelligent applications. Determining a user's posture on furniture generates information that can be used in various areas, such as home automation or Ambient Assisted Living. For example when a user lies on a bed, the system can deduct that he will remain there for a while, causing lighting and heating in adjacent rooms to be adjusted. If the user is suddenly sitting at night, the system may anticipate that the user is going to the restroom, thus activating dimmed lights to prevent tripping and falling.

In this chapter, we aim at classifying the posture of a person using capacitive proximity sensors that are embedded into the environment. Classification refers to making a discrete observation, e.g 'a person is lying on a couch', derived from a set of incoming sensor readings. In this work we will present an approach based on capacitive proximity sensors that can be unobtrusively integrated into existing furniture, while providing reliable information about the physical activities of a subject.

There are numerous elements within a user's environment that can be used for sensing physical activities. For example, Beetz et al [BJK*07] have equipped a kitchen with several types of sensors, realizing various scenarios based on a robotic assistant for supporting activities of daily life. One can also find cooking- and food-related systems like a diet-aware dining table [CLC*06], which recognizes different dishes (using RFID tags) and their weight (using a pressure-sensing surface). Similarly specialized is the eLab bench, which offers support to biologists in their daily lab work [ATB11]. Kivikunnas et al [KSK*10] have equipped a couch with capacitive proximity sensors for future application in posture recognition. Besides residential use-cases, detecting work activity has found widespread use in call centers and other fully computerized work environments to monitor productivity. In classical office work settings, activities carried out without computers are still commonplace [JS09], making computer-only work tracking systems insufficient for widespread workplace usage.

6.2.1. Classification Approach

By using capacitive proximity sensors, we can measure the proximity of a person's body. Based on fusing multiple sensors, it is possible to recognize patterns in the measured data. The sensors deliver continuous signals, which are sampled with a low frequency, e.g. 10 Hz, and normalized to an interval between 0 and 1. In the next step, overlapping short-time windows (e.g. with a length of 1 second), containing samples from all sensors, are built. Next, we extract relevant information for classification, the corresponding features, from these short-time windows. Typical features for user posture classification are the empirical mean and the standard deviation of a short-time window. For example, we may extract the empirical mean from each sensor and use it for classification. We aim to recognize a discrete class from the extracted feature vectors. A class may be represented by a statement like '*one person is sitting at the right side of the couch*' that reflects the user's situational context.

To cope with the complex task of classification, we need to learn from experience, making use of an annotated training set of feature vectors and the corresponding classes. The main goal of classification is to reliably identify a class for unknown feature vectors. Thus, generalization is a very important property of a classifier. In the following, we employ Naive Bayes, decision trees, and RBF networks as classification method.

6.2.2. The Smart Couch

We have chosen a couch in a living room as a subject of our investigations. The generated context can be used for energy saving purposes, e.g. shutting off lighting in other rooms. The data can also be the basis for controlling ambient parameters like lighting and multimedia equipment. Especially lighting may react differently to sitting and lying persons. We have equipped an ordinary couch with eight capacitive proximity sensors and applied various classification techniques to the generated data. We evaluated different classifiers by testing the prototype system on a diverse group of persons. Results show that such a system is reliably able to recognize various user postures, even if body mass and height differ strongly. Moreover, interviews performed with our test persons strongly indicate that the unobtrusive nature of capacitive proximity sensors will increase the user acceptance in actual applications, compared to camera-based systems. The interviews revealed that, using camera-based systems, people feel particularly observed and consider popular recent data leaks.

Since there are numerous scenarios for applying user posture classification, we have developed a generic framework called *SenseKit* for posture classification tasks. SenseKit is an integral part of OpenCapSense, our capacitive prototyping system presented earlier [GPBB*13]. Different scenarios, e.g. posture-detecting chairs or beds, require different types and numbers of capacitive sensors, in order to reliable detect the required postures. Sensekit tackles these classification tasks and provides additional functionality, most notably visualization and evaluation of the processing pipeline. It is based on a configurable dependency injection framework that allows all components (classifiers, sensors, feature extractors, etc.) to be dynamically combined.

Apart from SenseKit's classification and digital signal processing abilities it also implements training and visualization components. Sensor readings, as well as the final classification results, may be visualized and presented in an effective way. We have integrated various machine learning algorithms into our framework. Most algorithms are adapted from the WEKA Machine Learning Project [Uni11]. Moreover, we integrated WEKA's explorer into our framework, a comprehensive tool that provides functionality to evaluate recorded training data.

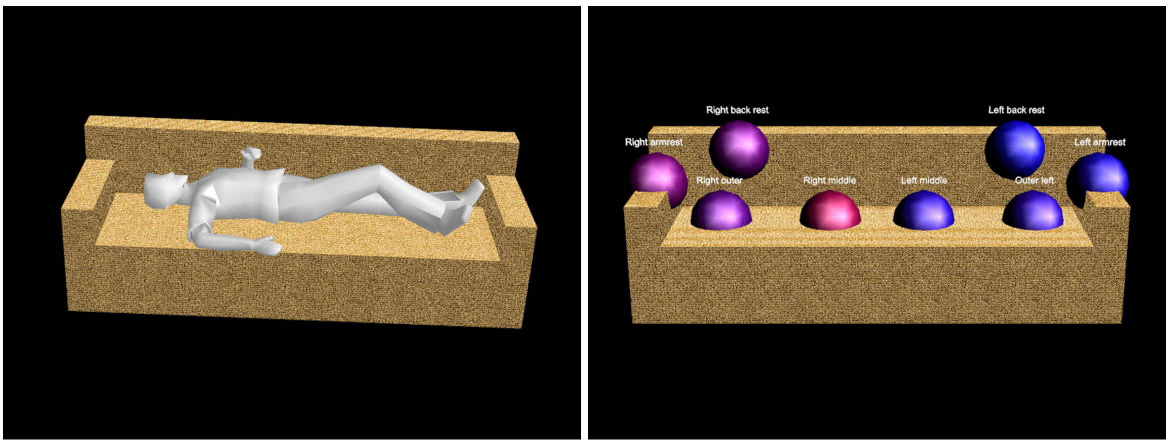

Figure 6.16.: Left: Visualization of a classification. Right: Visualization of sensor readings

Our prototype is an ordinary couch augmented with capacitive proximity sensors hidden underneath the upholstery. In order to prove our methodology, we intended to classify various sitting and lying positions (6.16 at the left) for one and two persons. We deemed eight sensors to be sufficient in establishing a good data basis for classification of the various postures. The final design consists of two sensors, placed underneath both armrests, two sensors in the back rests and another four sensors underneath the sitting area, as shown in figure 6.16 at the right. Even though the sensors are up to 15 centimeters away from the user and covered by upholstery and wood (see 6.17), we are still able to retrieve good measurements of body mass proximity. This is supported by the fact that some electrodes have pressure applied to them, causing a geometric deformation that also affects the output signal.

Figure 6.17.: An ordinary couch has been equipped with capacitive proximity sensors that have been set up under the upholstery and wooden elements

In order to determine a suitable test set, we are distinguishing nine different possible postures on a couch that can be performed by one or two persons. The 18 test persons were given simple written instructions to perform the desired postures. The persons performed all postures, relaxed and without restrictions in their movements, for approx. 30 seconds. Similar postures were always interrupted by unrelated postures. The data set contains a training set, with data from 9 test persons, and a test set with data of another 9 persons. Both, the training and test data set, were recorded on different days with different test persons. Additionally, we recorded body weights and sizes (figure 6.18), since those are the main properties affecting sensor measurements. Our training set consists of 2829 instances (about 24 minutes), whereas our test set consists of 2312 instances (about 20 minutes).

In the regarded scenario, we use the empirical means of our eight sensors, which are extracted from a short-time window, as feature vector. Overlapping short-time windows are passed to the classifier every second, containing the last 2 seconds of sensor readings.

Three classifiers were evaluated on our data set. We evaluate the performance of the Naïve Bayes classifier, decision trees and RBF networks. To measure the performance, we consider the metrics of precision and recall. As each sensor has individual characteristics, e.g. caused by different electrode sizes, the evaluation results are not symmetrical concerning the geometry of the couch. Furthermore, each sensor produces an individual amount of noise that has to be taken into account. The performances of the three classifiers are shown in table 6.19.

Our evaluation results show that the Naïve Bayes model does deliver inferior results compared to more sophisticated models, such as RBF networks. his fact is mainly caused by the very strong assumption of conditional independence, which is not satisfied in user posture classification scenarios. However, the Naïve Bayes model provides sufficiently precise results, as well as efficient training and data analysis. We retrieve an overall recall of 92.2% and a precision of 90.6%.

The evaluation of decision trees, built with the C4.5 algorithm, shows similar results as the Naïve

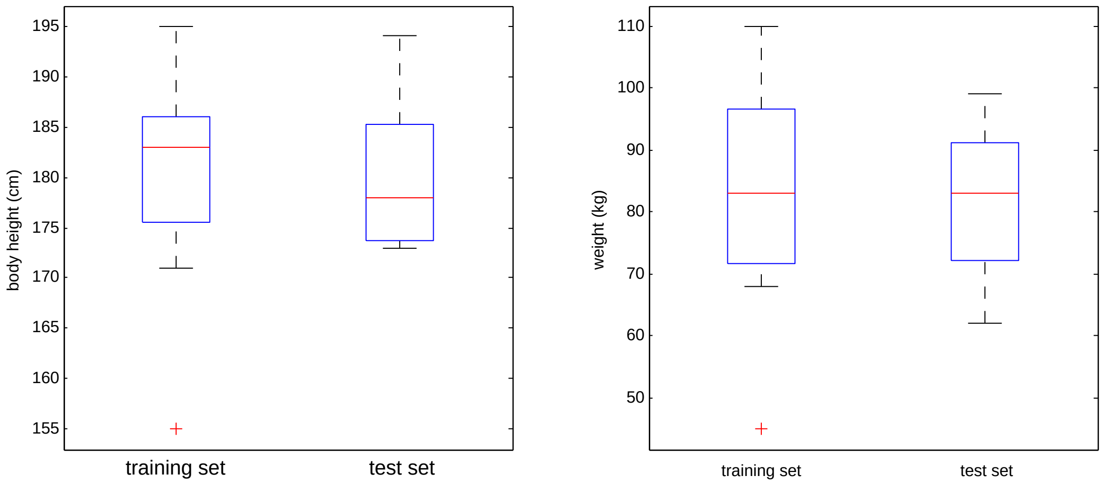

Figure 6.18.: A box plot of body heights and weights in our data set. The blue box denotes data from lower to upper quartile, the red dash denotes the median and red crosses mark outliers.

	Naïve Bayes		Decision Trees		RBF network	
Class	Prec	Rec	Prec	Rec	Prec	Rec
sitting outer left one person : OL	0.92	0.97	1.0	0.84	1.0	0.99
sitting middle left one person : ML	0.99	0.60	0.99	0.90	0.98	0.88
sitting outer right one person : OR	1.0	0.78	1.0	0.63	1.0	0.96
sitting middle right one person : MR	0.96	0.93	0.93	1.0	1.0	0.95
lying head right one person : LR	0.77	1.0	1.0	0.89	0.92	1.0
lying head left one person : LL	0.85	1.0	0.7	0.95	0.98	0.99
two persons sitting together : TT	0.77	1.0	0.95	1.0	0.87	1.0
two persons sitting gap : TG	1.0	0.99	0.98	0.64	1.0	0.99
no person : NP	1.0	0.92	0.66	1.00	1.0	1.0
Weighted average	0.92	0.91	0.91	0.87	0.98	0.97

Figure 6.19.: Evaluation results for the three classifiers

Bayes model. Classes like '*two persons sitting together with a gap*' were sometimes classified as lying postures. Moreover, many activities with lower sensor measurements, e.g. caused by a low body weight, have been classified as '*no person*', resulting in a poor precision of 66.4% for this particular class. The overall recall is 87.3%, whereas the overall precision is 90.7%.

Table 6.20 shows the confusion matrix of an RBF network, evaluated on our test set. We can see that some sitting postures on the right of the couch have been classified as *lying head right* postures, leading to a lower precision for this class. Furthermore, we retrieve a low recall for *sitting middle left* postures, as they are often misclassified as *sitting together* postures. In general, RBF networks perform very well on the test set with an overall recall of 97.5% and a precision of 97.2%. The determined clusters and their corresponding weights indicate that all sensors contribute equally to classification.

We can conclude that RBF networks are a robust classifier model with a high accuracy for user posture classification in our scenario. The generalization abilities of this classifier are coping well with the variation of body heights and weight of the different test persons. We have identified a decent generalization ability as an essential requirement for classifiers in user posture classification.

We have shown that capacitive proximity sensors are well-suited to give robust and reliable infor-

	OL	ML	OR	MR	LR	LL	TT	TG	NP	Prec	Rec
sitting outer left one person : OL	296	1	0	0	0	2	1	0	0	0.99	0.987
sitting middle left one person : ML	3	227	0	0	0	0	28	0	0	0.983	0.88
sitting outer right one person : OR	0	0	253	0	11	0	0	0	0	1.0	0.958
sitting middle right one person : MR	0	0	0	243	12	0	0	0	0	1.0	0.953
lying head right one person : LR	0	0	0	0	260	0	0	0	0	0.919	1.0
lying head left one person : LL	0	3	0	0	0	254	0	0	0	0.981	0.988
two persons sitting together : TT	0	0	0	0	0	0	197	0	0	0.872	1.0
two persons sitting gap : TG	0	0	0	0	0	3	0	212	0	1.0	0.986
no person : NP	1	0	0	0	0	0	0	0	306	1.0	1.0

Figure 6.20.: Confusion matrix for the RBF network classifier

mation about a user's context, proven in an evaluation with 18 different test persons of diverse body height and weight. Using only eight sensors in our couch example we have achieved a reliability of more than 97% in eight different postures using RBF network based classifiers. The classification based on machine-learning methods is easily implemented, trained and can be visualized by using the created SenseKit framework. We have achieved a fine-grained and reliable detection of user application context that can be used by intelligent systems to control the environment.

Open issues are the reliable detection of nearly similar postures, e.g. one person lying and two persons sitting. Most issues related to this topic can be solved by simply using more sensors. Even though SenseKit is supporting this the higher costs and complexity of the used hardware are undesirable. We intend to test other physical sensor configurations that could achieve better results as our current prototype. However, a well-defined theory and methodology, that describes the ideal distribution of sensors within the furniture is highly desired. Given the nature of capacitive proximity sensors and the highly complex distribution of electric fields, another option would be to use a simplified model to simulate the sensor values within the furniture and apply optimization strategies to achieve a good sensor configuration.

6.2.3. Smart Working Surfaces

The demographic change in many industrialized countries and consequential restructuring of social security systems leads to an increased number of elderly and disabled people in working life. Considering modern societies, computer work and activities that require a seated posture are one of the major risks for an employee's health, often resulting in a lack of exercise and stress to the spine.

These risks can be partially avoided by an ergonomic workplace that offers a personalized technical assistance in suitable situations. For example, additional lights can be switched on when the employee starts an activity that is related to reading documents. Moreover, the height of the table or parameters of the chair can be automatically adjusted to working situations. An ergonomic workplace utilizing assistive technology can not only help people avoid developing health issues, but can also aid people with pre-exisiting issues have a less distracting and more productive work experience. On the following pages, we present a method for recognition of working situations based on an array of capacitive sensors placed under a desk's wooden surface. Based on this work, it is possible to realize assistive services in the workplace to make it more ergonomic, efficient and supportive in the most relevant everyday desk work tasks.

The system used in this work is deployed under the surface of an ordinary desk, as illustrated in Figure 6.21. The grid is composed of three conductors placed horizontally and five conductors placed vertically. Using this setup, we can detect objects like hands and body parts located 10 cm above the desk, with a sensor update rate of approximately 50 Hz. The setup is the basis for subsequent feature extraction and

Figure 6.21.: The smart desk is equipped with a 3 by 5 grid of capacitive proximity sensors. The sensors measure the proximity to a user's body parts, for example the knees placed below the table, or the hands placed upon the table.

classification. Similar to the previous section, we employed loading-mode sensing, which measures the capacitance between an electrode and its surrounding environment.

We used the OpenCapSense evaluation toolkit, which is suitable for rapidly prototyping capacitive proximity sensing applications [GPBB*13]. Therefore, we placed eight loading-mode sensors under the tabletop surface and connected them to the wires under the desk's surface. The loading mode sensors were then attached to the OpenCapSense board by using standard USB cables. We used OpenCapSense's measurement and evaluation application Sensekit to record activities for later evaluation with the WEKA machine learning framework[1].

As a first processing step for recognizing working situations, we extracted time-windows of five second length from the eight proximity sensors. We then calculated the mean and standard deviation for each sensor window. Moreover, we extracted the center of mean and the center of standard deviation from all sensors. This can be achieved by weighting the sensors' x- or y-positions with the corresponding mean or standard deviation. These features acted as an input vector for later classification of the performed activity. The sensors were configured with a high update rate of 50 Hz enabling us to capture fast movements. Using the features extracted from the time-window, we applied RBF Networks for classification.

In our experiment, we were able to show that our approach is a promising concept for classifying working situations among different users. We also investigated if the classification approach can be generalized for all users, enabling cross-user classification without separately annotating training data for each person. However, we expected that the working situations are often carried out very differently, highly depending on the specific person and habits. Figure 6.22 shows an exemplary measurement result from a working situation. We can see that the sensor values reflect the placement of the user's hands and the proximity to his knees.

We identified the following classes that are typical for the presented office scenario: typing on a computer (employing only the keyboard), mousework (employing only the mouse), reading a book, phoning, pause, hand-writing and no person. Figure 6.24 shows two exemplary activities, that were

[1] http://www.cs.waikato.ac.nz/ml/weka/

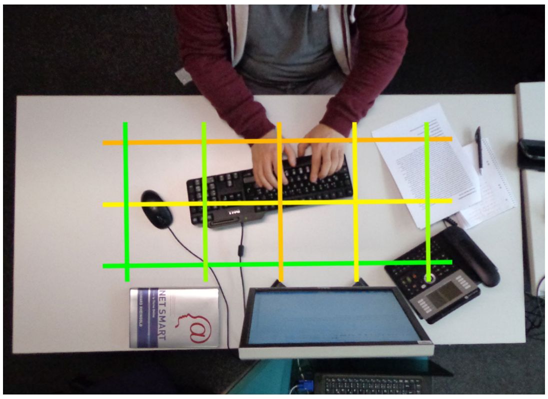

Figure 6.22.: Exemplary visualization of sensor values, depending on an activity. High sensor values are marked in red, low sensor values are marked in green.

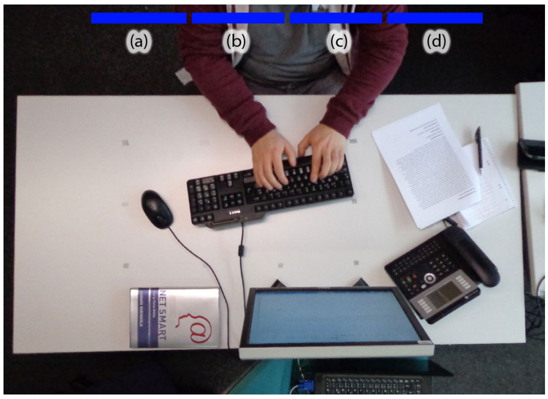

Figure 6.23.: The office chair's positions were split into five discrete classes: (a) outer right, (b) middle right, (c) middle left, (d) outer left and no person.

carried out in our evaluation. In addition to these common working activities, we aimed to recognize the office chair's position, illustrated in Figure 6.23. This position was classified separately employing five classes: outer right, middle right, middle left, outer left and no person. Even though the person might not always place the hands close to the table, it is possible to detect the proximity of the knees to make inferrations about the office chair's position.

Our test set consists of the activities of 12 persons who carried out each activity for approximately two minutes. The activities were always interrupted by non-related activities, such as moving away from the table or walking. In order to evaluate scenarios like reading or typing, we placed several office-related items like keyboards, mice and books on the table. These items have a very low impact on the sensor values, as they are not grounded and only slightly influence the measurements with their permittivity. The main goal of the evaluation was to identify working situations on the given data basis. Furthermore, we aimed to find out if a single training set can be shared among all participants and if the working situations of unknown particpants can be reliably classified. In order to evaluate the possiblity of having a shared data set, we performed a 4-fold cross validation on the recorded and annotated data of all participants.

With the RBF network classifier, we achieved an overall accuracy of 93.2 % for the four different desk chair positions. Splitting the test set into six participants for training and six participants for testing, we achieved an overall accuracy of 70.5 %. The reason for this lack of precision can be assumed to lie in the great variety of sitting postures and the different ways of placing one's arms on a desk's surface. Regarding the confusion matrix, it is obvious that office chair positions are often misclassified in their neighboring ones, such that misclassifications will not have a very negative effect on later applications using this data.

Considering the seven different working situations, we achieved an overall accuracy of 81.8 % for a 4-fold cross-validation. As the hand positions are very similar for reading and writing activities, these classes were often confused. With an accuracy of 93.7%, the class *mousework* showed the best performance for the given test set. The *pause* activity was classified with a very poor accuracy of 59.9 %. Regarding this activity, there were many variations in the body posture (for example leaning back) and the type of activity (e.g. eating chocalate).

When splitting the data set into a dedicated test and training set of six participants each, the classifier

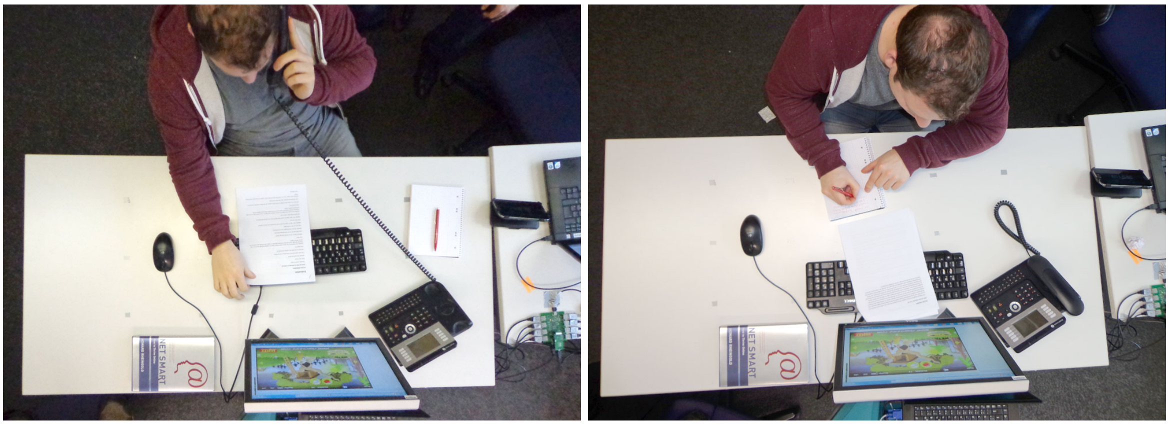

Figure 6.24.: Two exemplary activities, carried out above the smart desk: phoning and writing. The position of both hands is very different and can be exploited to distinguish between the two activities.

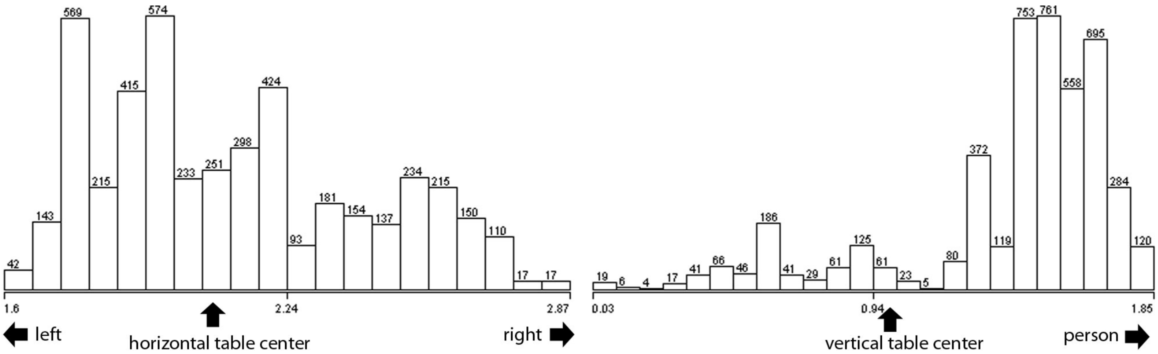

Figure 6.25.: The two center-of-mean features in x- (horizontal) and y-direction (vertical). The number above the bars represent the number of instances in the data set.

could achieve an overall accuracy of 49.8 %. Therefore, we must conclude that office activities are highly individual and must be trained in advance with each person. Moreover, the placement of the electrodes is not optimal, as Figure 6.25 shows. The plot for the y-axis center-of-mean shows that most activity was performed in the first half of the table, the area which is close to the person. Thus, it would be reasonable to deploy more electrodes in this area to achieve a higher resolution. The x-axis center-of-mean reveals that the placement of tools, such as a phone, and the user's characteristics, such as being right-handed, leads to a very unbalanced usage of the two tabletop halfs.

Using self-capacitance measurements from various electrodes enables to gather information about the working situation with a minimal amount of required hardware. We have created a prototype system based on the OpenCapSense rapid prototyping toolkit [GPBB*13] and performed an evaluation with 12 users. We tried to differentiate six different working situations associated to a typical office employment (typing, mouse-work, reading, hand-writing, pausing and talking on the phone). The results have shown that theses tasks are varying strongly between the different persons and it is difficult to correlate training data from one user to measurements of another. We can therefore conclude that a single array of sensors is not sufficient to reliably detect working situations. However they form a solid base for this approach and can be easily combined with other systems.

6.3. Summary

In this chapter, I presented various techniques that can be used to derive the situational context of a user. Here, methods for recognizing physical activities represent an important concept that contributes to this goal. Due to the emerging trend to smartwatches and wearable computing, it is foreseeable that additional sensing modalities will be required in the near future. In this field, I applied capacitive sensors to measure the proximity and nature of nearby objects. Based on the measurements, it becomes easier to recognize activities of daily living, such as drinking. Especially for such activities, capacitive sensing reveals its advantages, as the approach to objects, such as tables and the user's mouth can be quantified. Besides implicit interaction, the sensors can also be used very elegantly for explicit interaction. Using multiple capacitive sensors would enable wristband-based interaction, counteracting the problem that the interaction area on smartwatch displays is rather small.

In the field of assistive technologies, stationary deployments in furniture can play a key role in the future. By integrating information about furniture usage into the context knowledge of a system, the environment is able to respond intelligently to recognized patterns. I described two use-cases based on capacitive proximity sensing: recognizing working situations on tabletops and recognizing user-postures on a couch. Based on the gathered data, it is possible to detect persons reading at a tabletop, and react by automatically increasing the light level.

Especially due to their unobtrusiveness and invisible placements, the use of capacitive proximity sensing can contribute significantly to implicit interaction. The modality's ability to measure proximity allows for placements underneath non-deformable and pressure insensitive surfaces, such as tabletops or wooden parts underneath upholstery. In wearable computing, the low energy consumption and the low hardware requirements are a very important advantage of capacitive sensing compared to other technologies.

7. On the Disappearance of Affordances

In implicit interaction, the common goal is to contribute to an understanding of users and their situations [Sch00]. Especially in this domain, technology tends to disappear and becomes an invisible part of a user's environment. Although to a smaller extent, the same applies for explicit interaction. For example, the opportunity of controlling a TV using gestures is not directly apparent. Only by experimenting or by searching for a tiny camera placed in front of the TV, it is possible to derive such interaction opportunities. This induces problems in all generations, especially those who are not used to the fact that technology may be embedded in everyday objects or within the environment. The same yields for younger people, although they are usually able to adapt faster to new technologies. In this chapter I address this problem with respect to a capacitive gesture recognition system (Figure 7.1).

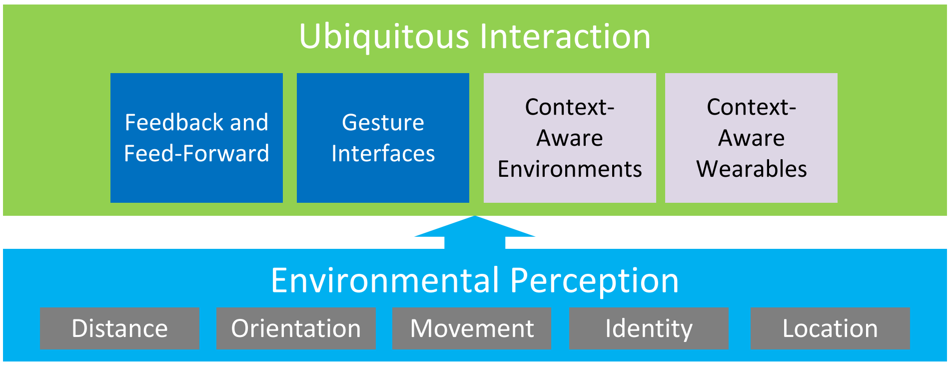

Figure 7.1.: In the domain of explicit interaction, I investigate a gesture interface which is able to provide feedback and feed-forward clues to the user.

The central problem I am approaching can be summarized very shortly: the affordances of many interaction systems tend to disappear. In his very inspiring book *The Design of Everyday Things*, Don Norman describes the term *affordance* as a signaling mechanism to humans [Nor02]. For example, a button signals a user to push it, and ideally, it also reveals the outcome of that action. In contrast to a *signifier*, which gives an indication *where* interaction can take place, an affordance conveys *what* action is possible. Personally, I can imagine that many people have experienced a situation in which the affordances deceived [Nor02]:

> *"I have seen people trip and fall when they attempted to push open a door that worked automatically, the door opening inward just as they attempted to push against it."*
>
> — Don Norman, The Design of Everyday Things [Nor02]

The problem occurs when technology becomes invisible, for example when technology moves from simple push or rotate buttons to gesture recognition applications. This leads to confusion and requires training and instruction on the different types of gestures and their outcomes beforehand. Unfortunately, this is exactly the opposite the research community aims to achieve. Gesture recognition can be associated to the category of techniques that should make interaction more natural. But currently, we are missing out real life and thus those persons who are not familiar with these systems. In our Ambient Intelligence lab, we have a number of appliances that are built upon gesture recognition, let them be either camera-based or capacitive. Unfortunately, all systems including mine are not immediately usable by novice users.

This made me struggle over the question which kinds of affordances and signifiers are still necessary when an intelligent environment could perfectly recognize natural interactions. I believe that a large amount of signifiers could vanish in that case, but the majority of affordances will remain. Don Norman introduces the following factors that make up a good product design [Nor02]. *Natural Mappings* take into account that an intention can be naturally mapped to an action, which causes a system's reaction. This is a rather big problem in gesture recognition systems, as they are only able to cover a subset of possible movements. For example, in my capacitive gesture recognition systems, people tend to move their fingers although sensing resolution is too coarse to capture these actions. *Design for Error* is very straight-forward: Errors can occur at any time and users will make mistakes. It is not the user's fault when his or her action are interpreted wrongly, rather the designer and thus me is responsible. *Providing feedback with multiple senses* is also important to give the user an impression of any action's results. Feedback can be deployed with different senses, such as haptically through touch, acoustically or visually.

Figure 7.2.: Gesture-recognizing surfaces based on capacitive sensing can act as a low-cost interface for interacting with entertainment system (left) or with smart doors, for example in public restrooms (right).

I experienced the lack of feedback as an integral problem in capacitive gesture recognition systems. Users are often not aware of their recognition boundaries and the actions they are able to perform. These two points induce the need for two approaches: feedback and feed-forward techniques. On the one hand, I show users how to use the system based on feed-forward information. On the other hand, users retrieve feedback on the interaction status by simple feedback techniques. In order to support more natural forms of mappings, I applied a visual feedback directly on the interaction surface. Figure 7.2 shows two examples, in which gesture-recognition can be enhanced with visual feedback. For example, it is possible to project a glowing shadow onto the device's surface to indicate that a hand has been recognized. In combination with gesture-controllable door, it is possible to indicate possible interaction patterns by natural mappings with colors and animations.

This chapter is based on two papers [GPBW*14b, GPBW14a] and includes results from Sebastian Beck's Bachelor thesis [Bec13]. The use of 'we' refers to the papers' authors Tobias Grosse-Puppendahl, Sebastian Beck, Daniel Wilbers, Steeven Zeiss, Julian von Wilmsdorff, and Arjan Kuijper.

7.1. Gesture-Recognition and Affordances

The success of gesture-based interaction in the context of entertainment systems initiated an increasing trend towards natural interaction paradigms in intelligent environments. Gesture recognition systems in smart environments represent an intuitive and easy way of interaction with the surrounding, for example allowing the user to control everyday devices using simple gestures. There are several technologies that are suitable to act as an input modality for gesture recognition. Commercial camera-based systems, such as the Microsoft Kinect [Pre10], are able to capture gestures and movements within a room. Other approaches employ the environmental noise in an environment [CGL*12] or use mobile phones for gesture recognition [BBRS06].

Stationary installed capacitive sensors can act as both, touch-sensitive and proximity-sensitive gesture-recognizing input modalities. Sensing proximity is especially suitable for recognizing user interactions within a well-defined interaction space. Considering scenarios that require gesture recognition without a graphical user interface, it is challenging to provide a suitable and meaningful feedback on the current interaction state to the user. This feedback can be provided by different modalities, for example visually or acoustically. Majewski et al. use a laser spot that visualizes a user's pointing direction perceived by the environment to disambiguate device selection [MBMK13]. When a device is selected, the spot delivers additional feedback on the successful selection by blinking. The authors of [SBW12] project a visual feedback directly on the user's body. The presented system provides hints on recommended hand movements and delivers feedback on the movements performed. Moreover, it is possible to use floating air as feedback, enabling accuracies of 8.5 cm at 1m distance [SPGI13]. Using electrodes attached to the body, Lopes et al. provide muscle-propelled feedback on the human body [LB13]. The approach is based on conducting currents between two electrodes, very similar to fitness appliances for enhancing body-shape. To our knowledge, capacitive proximity sensing devices have not been directly augmented with visual techniques to give feedback on the current interaction status and indicate possible interaction paradigms.

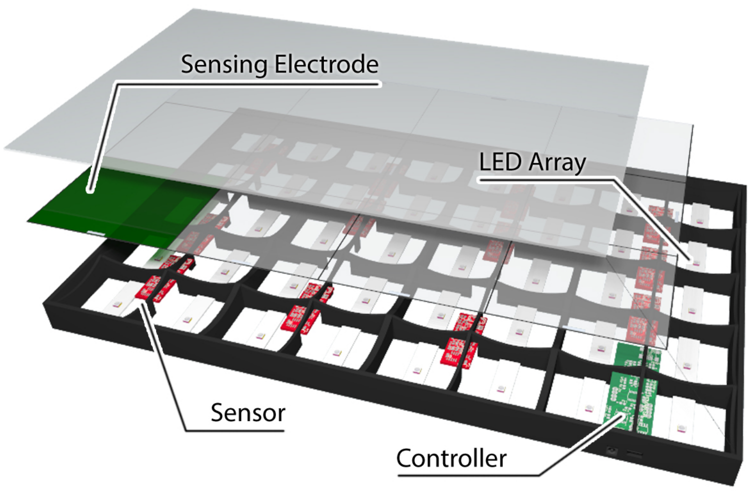

Figure 7.3.: The device consists of four main components: the sensors, the shielded electrodes made of transparent ITO, an LED array and a controller board. All components are interconnected by an I2C bus.

7.2. Hardware & Processing Chain

A conceptual drawing of Rainbowfish is depicted in Figure 7.3. It employs 12 sensors which measure the capacitance between the sensor's electrode and its surroundings, also known as a loading-mode measurement [SGB99]. The sensing electrode's surface builds up an electric field to any object in its surrounding. When a human hand approaches the sensing electrode, the capacitance increases. This effect allows for determining an approximate hand distance based on each sensor's measurement. By combining measurements of all 12 sensors, the hand's position can be inferred. In order to conduct the measurement, the resulting capacitor between the electrode and the surrounding objects is charged and discharged with a frequency of 500 KHz. The sensing electrodes are transparent PET foils with a conductive layer of Indium-Tin-Oxide (ITO), a material widely used in modern capacitive touchscreens. They have a size of 10 cm by 8 cm and consist of two layers: a sensing and a shielding layer. The shielding layer is necessary to avoid electronic interferences with the underlying hardware, such as the sensors and the LED array. All components are embedded into a 3D-printed grid structure. The electrodes are adhered underneath the device's top surface - a 3 mm thick layer of semi-transparent Plexiglas.

The device has a central Inter-Integrated-Circuit (I2C) bus used for interconnecting the sensors with a controller board. This controller board is responsible for scheduling the sensors and controlling the LED array. The measurements are performed concurrently to achieve a suitable temporal and spatial performance. However, when sensor electrodes are located side-by-side, a parallel measurement would affect the neighboring sensor. Therefore, in each measurement step only three sensors are activated in parallel to avoid interference. Using this method, we currently obtain 20 measurements per second for each sensor. Rainbowfish's controller sends the sensor values to a PC through a USB connection. The PC executes additional processing steps, like drift compensation and normalization, and determines the position of a user's hand. Based on this information, an application is able to send information on its execution state or supported gestures back to Rainbowfish. Instead of sending pixel-related data with a high update rate, the application sends lightweight function calls to trigger pre-defined visualization profiles. These profiles include illuminating the whole surface in a specified color and for a certain time, animating a swiping gesture, drawing colored rectangles, and glow effects on continuous 2D coordinates.

7.3. User Study

The overall goal of the user study we conducted was to explore the applicability of visual feedback on a gesture-recognizing surface in smart environments. We stated a number of hypothesis, which were investigated in the study: (*H1*) visual feedback increases the interaction speed, (*H2*) a novice user is able to handle an unknown system more easily, (*H3*) a user is able to recognize usage and system errors immediately, and (*H4*) the perception of visual feedback depends on the familiarity with the system. Therefore, we conducted a user study with 18 participants. The study consisted of four main parts: (1) which gestures would a user perform to trigger a certain action, (2) imitate gestures based on visual feed-forward animations, (3) interpret visual feedback and (4) use Rainbowfish in two exemplary applications for smart environments.

7.3.1. Perception of Feedback and Feed-Forward Visualizations

In order to investigate if users are able to handle a system more easily with visual feedback and feed-forward clues (*H2*), we conducted an experiment consisting of two parts. First, the participants were asked to perform certain gestures to reach an application specific goal without any animations shown on

Rainbowfish's service - for example by giving instructions like '*raise the volume of a media player*' or '*switch the light off*'.

The variety in which the gestures were carried out turned out to be very high and coherent. However, they showed substantial analogies to smartphone and tablet PC usage. Eventually some general statements can be made for certain goals, for example for instructions like *select an object*. This instruction mainly resulted in grabbing, tapping on the device's surface, or hovering over the object. When considering smart environments, it can be concluded that gesture-recognizing surfaces are very hard to handle without feed-forward information if only the functional goals are known to a user.

In the second experiment, the test persons were asked to imitate gestures based on feed-forward visualizations projected on the Rainbowfish's surface. Again, we investigated the variety of gestures which were carried out. The feed-forward animations are shown in Figure 7.4. Therefore, we exploited analogies to common touchscreen gestures (pinch-to-zoom, rotate, etc.), which led to a vast majority of correctly performed gestures (93.5 %). This supports the assumption that the presented feed-forward animations are a suitable way of representing the affordances of a gesture-recognizing surface.

In the following experiment, we presented each participant a number of feedback expressions displayed on Rainbowfish's surface. This experiment was conducted to explore how visual feedback provided by an application is perceived by a user (*H3*). Figure 7.5 shows a subset of feedback animations which were evaluated. As expected, a short green flash was associated with the acknowledgment of an action by almost all users. On the other hand, a red flash was associated to neglection or rejection. Yellow and blue flashes were mainly associated to a wide variety of meanings, such as *waiting* or *in progress*, which does not allow for any generalizable statement. Interestingly, more users were able to associate a green flash with a positive outcome when the complementary red flash was shown afterwards.

7.3.2. Evaluation of Applications in Smart Environments

In the next part of our user study, we investigated to exemplary applications which we developed for Rainbowfish. In the first application the participants controlled a home entertainment application - an image viewer - with gestures. This application consists of our gesture recognition device, as well as a screen for displaying the images. The second experiment solely employs the gesture recognition device without providing an additional graphical user interface. In this part of the evaluation, the users were asked to open, close and lock an automatic door by performing gestures.

7.3.2.1. Home Entertainment

In this experiment, an image viewer as an exemplary home entertainment application was evaluated. We placed our gesture recognition device in front of a screen that showed the image viewer application. In this setup, depicted in Figure 7.6, the user is able to manipulate the application's cursor by the position of her or his hands. The participant is able to scroll to both sides by placing a hand near the edges of the device. In the detail view, horizontal swipe gestures are employed to switch to the next or previous image. A vertical swipe gesture from top to bottom allows the user to return to the overview.

We implemented various types of visual feedback on the device. When a hand is recognized by the device, a blue glow effect follows the position of the user's hand, similar to a shadow. In the image viewer's overview the regions at both sides of the device are illuminated to visualize the possibility of scrolling (see Figure 7.7). When the hand remains above an image in the overview, the glow effect fades from blue to green to indicate a successful selection. At the time a gesture is performed, the device indicates the successful recognition by shortly lighting up in green (see Figure 7.8).

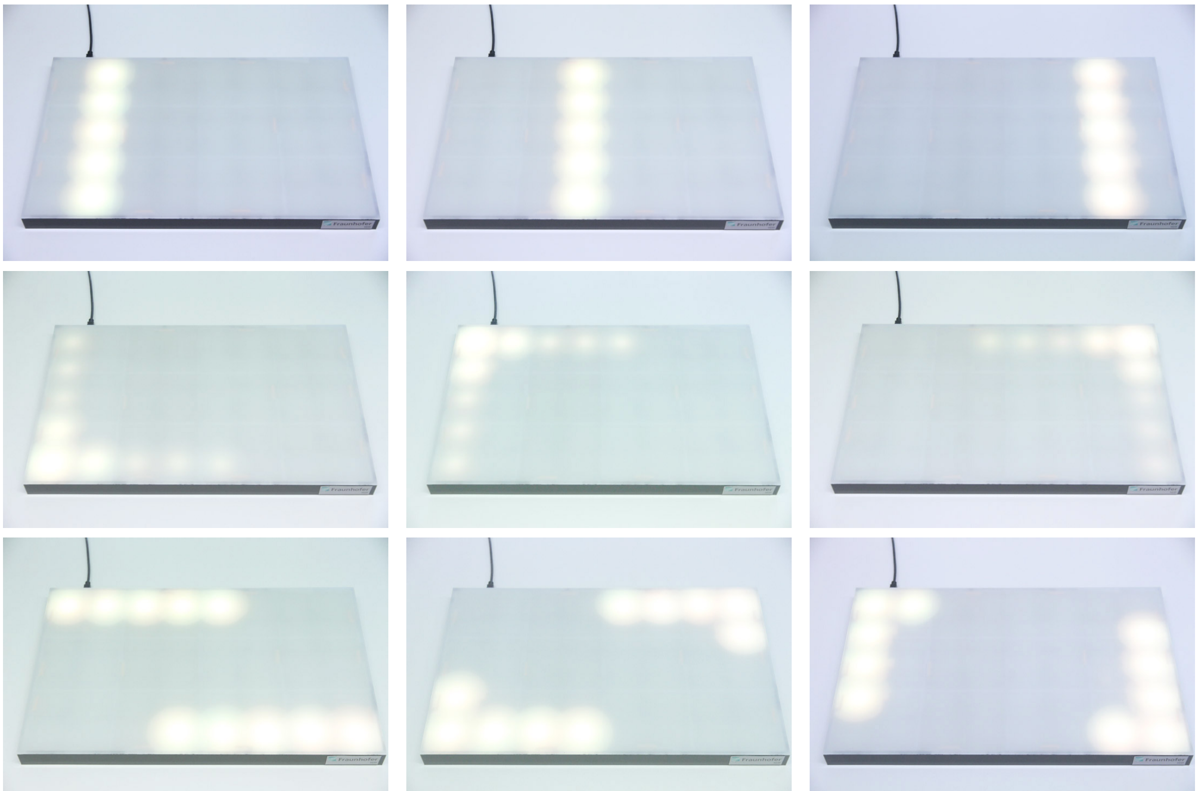

Figure 7.4.: feed-forward animations for gestures in front of Rainbowfish. The first column shows a swipe gesture from left to right, the second indicates a rotate gesture with a single hand, whereas the third visualization shows a two-handed rotate gesture.

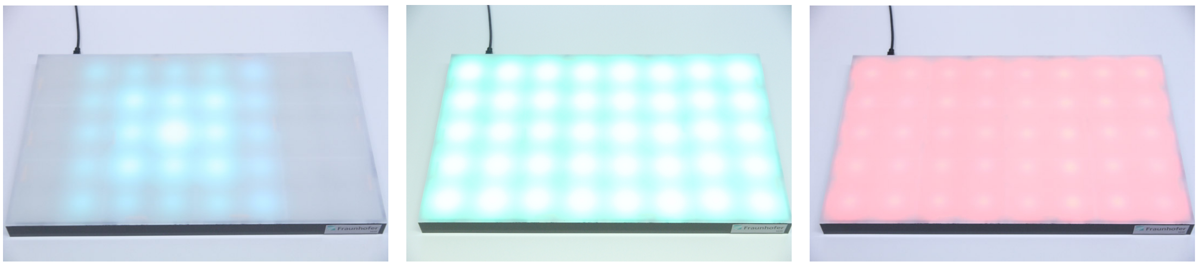

Figure 7.5.: Different types of feedback can be used to indicate certain application-specific outcome. In our study we asked the users to associate a meaning to the animations shown in the three images.

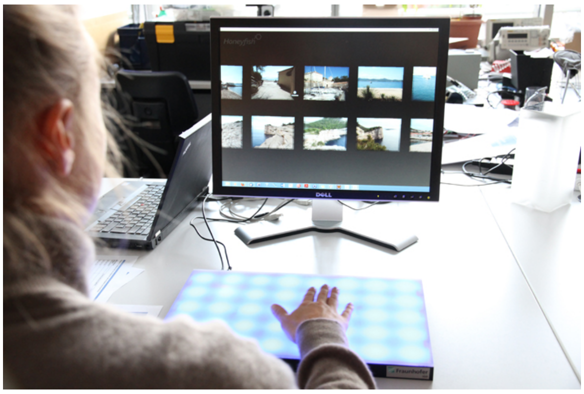

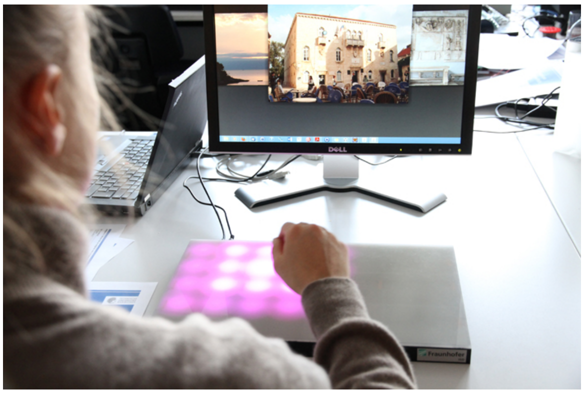

Figure 7.6.: In the image viewer application, a user is able to select and browse between images using gestures which are enriched with feed-forward animations and interactive feedback.

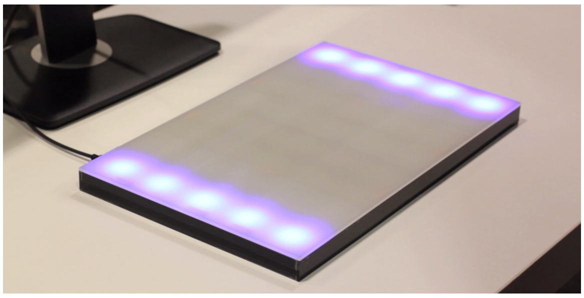

Figure 7.7.: Interactive regions are visualized with a glow effect. When the hand moves over the corresponding region, an application-specific action is triggered (e.g. scrolling).

Figure 7.8.: When a gesture is recognized successfully, the device lights up in green. Moreover, it is possible to indicate unrecognized or unsupported gestures by lighting up in red.

Every participant was instructed to perform a set of tasks, one group obtaining a visual feedback by the device and one without. In order to find out if visual feedback speeds up the interaction (*H1*) and makes usage or system errors visible faster (*H2*) we recorded the number of unsuccessfully recognized gestures and the resulting interaction speed by counting the number of actions in a given time span. Additionally, we asked qualitative questions on a Likert scale from 1 - 10 to investigate if the perception of visual feedback depends on the familiarity with a system (*H4*) and the interaction becomes easier for novice users (*H3*).

The participants were asked if they paid attention to the visual feedback provided by the Rainbowfish. One test person did not observe any feedback at all, because she was focused on the application shown on the television. Many other participants had a similar experience: they were not able to interpret the different effects and colors of the board because they focused on the application itself. Some could not associate their actions with a color or animation. Overall, the participants only showed a slight tendency to pay attention to the device's visual feedback (5.65/10 points) and supported them in their initial steps with the device (6.44/10 points). Despite the limited perception of visual feedback, the majority of users did not feel disturbed by the illuminated surface (3.00/10 points).

In conclusion, the evaluation of a home entertainment application showed that two visual feedback mechanisms - the graphical user interface and the gesture recognizing board itself - were not necessary

Figure 7.9.: The minimalistic feedback shows the state of the door - which is currently locked (red).

Figure 7.10.: The extended feedback also indicates when a hand approaches the gesture-recognizing surface.

for the majority of users. Nevertheless, novice users or users who experienced problems during the interaction benefited from the visual feedback and feed-forward animations (*H4*). An additional positive aspect can be seen in the influence of the Rainbowfish's multicolor lightning on the intrinsic motivation of a user. It was mentioned by many participants that they liked the device and especially the colors, and were motivated to start interacting with it (*H3*). The interaction speed could not be increased by providing feedback and feed-forward information (*H1*).

7.3.2.2. Contactless Door-Closing Mechanism

We also conducted an experiment on controlling parts of an intelligent environment without using a graphical user interface. Therefore, Rainbowfish may be incorporated into walls, doors, or home appliances like cooking plates. We built an automatic door that can be controlled using gestures - for example to be used in public restrooms. A user is able to lock, unlock, close and open the door by performing horizontal movements in front of the device. The device delivers interactive feedback on the interaction state and gestures that can be performed. The automatic door control has three possible states, with the related colors: open (green), closed (yellow), and locked (red).

We compared two different types of visual feedback. First, a minimalistic feedback is provided by illuminating the device with the color of the current door state (see Figure 7.9). Second, we also visualized the gestures that are required to switch to the next state (see Figure 7.10). For example we visualized a red swipe gesture within the 'closed' state of the door to indicate that the door can be locked. Therefore, the corresponding colors of all states were used to visualize the required gesture.

Rainbowfish's output was essential to recognize the state of the door, as the *closed* and *locked* state cannot be differentiated by the user. The participants acknowledged that they directly focused on the visual feedback (8.48/10 points), even if they were not novice users (*H4*). At the same time, the users felt slightly more disturbed by the visual feedback than in the first experiment (3.67/10 points). Nevertheless, most of the participants could interpret the correct meaning of color and animation correctly. However, the interaction speed did not improve (*H1*). The opinions about the two provided modes varied strongly among the participants. Some of them mentioned that it was not necessary to animate swipe gestures because of their convenience, and a simple state-dependent feedback was sufficient for this use-case. The majority of all participants experienced the animated feedback to be very helpful.

7.3.3. User Study Summary

It can be concluded that feed-forward animations and feedback can help novice users to control devices by gestures in a smart environment (supporting *H2*). Visual feedback and feed-forward information helps this group of users when experiencing usage problems (support *H3*). Users who are familiar to a system do not benefit substantially from feed-forward animations (supporting *H4*). Moreover, the visualizations on Rainbowfish's surface had no influence on the interaction speed (not supporting *H1*). When providing an additional graphical user interface, the perception of feedback and feed-forward animations is very limited. This supports the assumption that a system with visual feedback should be deployed as a stand-alone input modality within a smart environment.

Many users also criticized time delay as well as limited interaction distance. These problems are mainly related to technical issues, which resulted from the transparent electrode material. Mechanical deformations of ITO foil can lead to slight damages of the coating, and thus, a decreased conductivity. This effect resulted in several problems months after building the device. Furthermore, when the material is deformed due to mechanical influences (e.g. by a tap on the surface), the capacitance may change rapidly and lead to unexpected behavior. In the future we will strongly focus on more resilient materials, for example thin conductive layers of silver on PET foil. Also, we aim to achieve interaction distances of 30 cm increasing the voltage levels from 3.3 V to 12 V.

7.4. Summary

Rainbowfish is a contribution to the third research challenge on new interaction concepts. The system is capable of delivering interactive visual feedback and feed-forward information. Two demonstration applications were developed to evaluate different usability aspects in smart environments. The evaluation revealed a set of possible inferences for the usage of visual clues on gesture-recognizing surfaces.

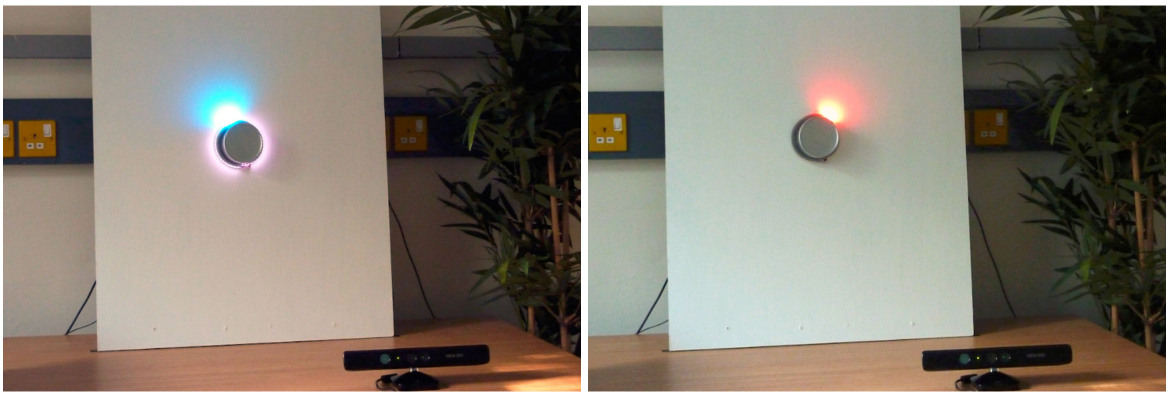

Figure 7.11.: The gesture thermostat projects light around the device indicating the heating level (feedback) and possible gestures to change it (feed-forward) [FBL14].
[FBL14] Reprinted by permission of the authors.

Feedback and feed-forward information is especially helpful for novice users who are not familiar with the corresponding gesture recognition system. When a graphical user interface is employed, experienced users often do not notice visual feedback provided on the gesture-recognizing surface. On the other hand, when no GUI is provided, visual feedback also helps experienced users to interact with the system. Having completed the experiments, participants looked forward to use our applications - resulting in a multitude of ideas where technology could be used in the future. Especially public sanitary installations, like toilet flushes, toilet doors or doors in general were mentioned. Besides that, a water tap with a gesture-controlled temperature and water regulation was the most popular idea. Moreover, many applications within a living environment were mentioned, especially in the kitchen and the bathrooms where hygienic requirements are needed. Situations in which the user has sticky hands or carries things can be simplified by gesture-recognizing fridge doors, cookers or drawers. Various other ideas included the control of ambient lightning, gaming, multimedia applications and interactive furniture. A concept heading towards a similar direction as the participants' ideas was recently presented in [FBL14]. As depicted in Figure 7.11, light projections around a thermostat convey feedback on its current status. Moreover, gestures are indicated by projecting clockwise or anti-clockwise movements around the device.

Rainbowfish is a step towards the goal of making affordances visible. The ultimate goal is to provide the same or even more signaling mechanisms to users as physical objects can transport. A requirement for achieving this vision is to convey feedback and feed-forward information with more than one modality. In the future, it is necessary to work with all human senses, which include haptic, acoustic, and visual clues. All sensations must be combined intelligently to provide natural mappings and usage constraints. Especially the latter argument will be very hard to achieve, as in-the-air gesture recognition is hard to constrain in terms of possible movements and actions.

8. Conclusions

8.1. Summary

With this work, I contribute towards the goal of making interaction with technology more natural, efficient, and enjoyable. This objective is achieved by retrieving information about a user's situation and actions. The more a device or an environment understands and perceives about the user, the more intelligent it is able to respond. The system can thus contribute to the users' goals, which are inherently revealed by their implicit and explicit interactions.

Having Mark Weiser's vision of UbiComp [Wei99] in mind, I decided on my research focus in capacitive coupling technology. The main driver for my research was the goal of lowering the design trade-offs when using this type of technology. Capacitive sensing offers many desirable properties, such as low cost and low energy consumption. Sensors can be embedded easily in the environment with unobtrusive placements. The conductive properties of a human body enable the design of human-centric, highly interactive systems. Unfortunately, developers are often not able to take advantage of these properties when the perceptional capabilities are not sufficient for the desired application. Therefore, I aimed at extending these capabilities to design novel systems for ubiquitous interaction.

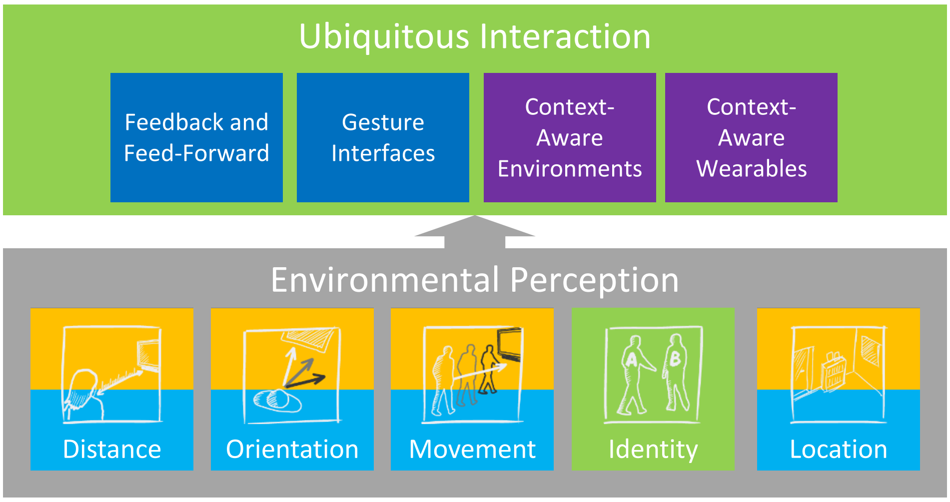

Figure 8.1.: A detailed perception of the environment [GMB*11] allows for ubiquitous interaction.
[GMB*11] © 2011 Association for Computing Machinery, Inc. Reprinted by permission

Figure 8.1 shows the conceptual outline of my thesis. It comprises two areas: environmental perception and ubiquitous interaction. Perceiving the environment is a necessary prerequisite for enabling ubiquitous interaction. Novel interaction opportunities can arise when this perception is extended to new areas. I applied proxemic interactions as a basis for environmental perception [GMB*11]. The concept of proxemics was first investigated by Hall, who defined human-centric proxemic zones. Transferring proxemics to UbiComp and HCI enables to model the spatial relationships between entities, such as users and devices. These relationships can be expressed in terms of five proxemic dimensions, which comprise distance, orientation, movement, identity, and location. Measuring these dimensions enables the creation of ubiquitous interfaces. Such interfaces can be characterized by technology disappearing in the environment, natural interaction metaphors, and collaborating devices.

In the field of *environmental perception*, I extended the capabilities of capacitive sensors within the proxemic interaction dimensions. I presented OpenCapSense, which provides means for prototyping UbiComp and HCI applications. It is targeted towards the dimensions of measuring distance, orientation, movement, and location. OpenCapSense is able to measure the proximity to body parts with a high spatial resolution at distances up to 50 cm. Although OpenCapSense covers four of five proxemic interaction dimensions, the dimension of identity is still missing. Measuring identity does not only refer to measuring a unique ID, it also incorporates measuring the inner state of an object. This information is not perceivable by external sensors and relies on mutual collaboration between objects. In order to cover this dimension, I introduced a novel methodology for capacitive near-field communication (CapNFC). CapNFC provides concepts for realizing communication among multiple objects based on capacitive coupling. Due to the conductive properties of a human, it can be regarded as a hybrid method for sensing and communication. It can be applied to communicate through air and the human body as well as measuring proximity. The inner state of an object is determined with auxiliary sensors, for example inertial measurement units. CapNFC and OpenCapSense provide the physical abilities of measuring all five proxemic dimensions. To infer detailed information about proximity, it is necessary to process the data intelligently. Here, Swiss-Cheese Extended comes into place - it enables to recognize object configuration in multiple degrees of freedom. The method allows for tracking multiple object trajectories in real-time, which can be classified to gestures or body movements. The three contributions complete the picture of proxemic interaction dimensions.

Based on the novel perceptional capabilities, I introduced new approaches for *ubiquitous interaction*. I differentiate between contributions in implicit interaction and explicit interaction. In implict interaction, the system unobtrusively recognizes the user's situational context. Deploying capacitive sensors within the environment allows for recognizing human activities. I presented two use-cases on posture-recognizing furniture: a couch and a smart desk. Based on the recognized user situation, it is possible to adjust the environment to the user's needs. In explicit interaction, a user triggers an action and awaits a response from the computing system. Providing natural interaction possibilities is desirable when explicitly triggering actions. Gesture-recognition systems represent an instance of such natural user interfaces. However, when no additional modalities are applied, they face the problem that affordances are invisible to the user. It is often not possible to directly interact with the system due to the disappearance of signalling mechanisms and the lack of natural mappings. I counteracted this problem by presenting a novel way to provide feedback and feed-forward information on gesture-recognizing surfaces. The surface is illuminated with LEDs and indicates possible gestural movements and reveals their outcomes.

A detailed perception of users and their environment is essential for natural interaction design. Retrieving fine-grained contextual information enables devices and environments to contribute intelligently to a user's goals and needs. Based on many advantages in terms of power consumption, deployment size, and unobtrusiveness, I decided to contribute to this goal with capacitive sensing. To make full use of the modality's desirable properties, my primary goal was the extension of perceptional capabilities. With my contributions I was able to introduce a variety of novel and innovative sensing and interaction approaches. This way, I could lower design trade-offs and move towards a better exploitation of capacitive sensing's advantages.

8.2. Scientific Contributions

I oriented myself on three different research challenges. In the following, I will outline my contributions to the scientific community for each challenge.

New capacitive sensing approaches: Chapters 3 and 4 contain contributions to the first challenge. In the first chapter, I present a rapid prototyping toolkit for capacitive sensing applications [GPBB*13]. OpenCapSense is the first prototyping platform that supports all three proximity sensing modes, introduced in Section 2.3.2. Compared to an existing prototyping toolkit [WKBS07a], it offers a significantly higher temporal and spatial resolution. The toolkit was validated with various prototypical applications which include a fall-recognizing carpet, an interactive art installation, and a gesture recognition system. As a second contribution to the first research challenge, Capacitive Near-Field Communication represents a method for mutual collaboration among objects [GPHW*14]. Previous works have primarily focused on specific use-cases, rather than conveying a bigger picture of capacitive communication. Here, my contribution is a novel conceptual basis which includes the identification of three different operating modes. The method was quantitatively evaluated based on a reference implementation. It was validated with three case studies that touch different areas of UbiComp and HCI. The two contributions provide the physical sensing opportunities for approaching the next research challenges.

Interpretation and fusion of data from capacitive sensors: My contribution to the second research challenge is discussed in Chapter 5. I presented Swiss-Cheese Extended, an object recognition algorithm based on elimination [GPBKK13]. Understanding how hands and or other body parts are situated is a vital information for context-aware systems. Previous approaches often relied on discrete classification or have rather limited continuous object recognition capabilities. Applying Swiss-Cheese Extended enables to recognize object configurations with multiple degrees of freedom in realtime. The method was validated with two case studies: a gesture recognition device and a device for interacting in front of a display. Qualitative and quantitative evaluations within the two case studies support the applicability of the presented method.

Interaction design based on novel perceptional capabilities: The first contributions to the third research challenge in the domain of implict interaction are presentend in Chapter 6. In order to raise the expressiveness of sensing in wrist-worn devices, I introduced a novel modality with capacitive proximity sensing [GPBB12]. An evaluation shows significant enhancements in recognizing activities when incorporating information about the proximity and nature of nearby objects. Besides wearable computing, stationary installed sensors in smart furniture can reveal information about human activity. With an augmented couch, I could reliably classify a high number of user postures with only eight capacitive sensors [GPMB11]. Furthermore, I outlined and evaluated a new approach for recognizing a user's activity on a desk [GPBK*13]. Chapter 7 describes a contribution in the field of explicit interaction. Rainbowfish helps users to interact with a gesture-recognizing surface by providing feedback and feedforward clues [GPBW*14b,GPBW14a]. The novel concept is based on projecting light on the surface of a capacitive gesture recognition system. Based on two exemplary applications, I evaluated the usability of the presented method and showed that users can benefit of this approach.

8.3. Future Work

My current research was mostly focused on capacitive coupling approaches. In the future, I would like to expand it to other modalities such as wireless signals [AKKM14,PGGP13], ultrasound [GMPT12], or acoustic sensing [OST13]. However, the goal remains the same: the unobtrusive perception of humans and their environment. The following ideas seem interesting to me and worth investigating.

Rethinking the touch-screen controller: In my previous work, I equipped *physical* objects with communication abilities [GPHW*14]. Transferring my ideas to the *virtual* world of a touchscreen would transform the touchscreen to a capacitive communication medium, as first shown in [DL01]. This would allow for realizing techniques like through-body transitions between surfaces [WB10] or pick-and-drop actions [Rek97, WB10]. While conventional touchscreens are tailored to sense touches, they are currently not good at recognizing physical objects on or near the screen's surface. In related work, various modalities like static magnetic fields, NFC, ultrasound or computer vision were used for identifying objects above the screen [CPL12, HIB*07, JGAK07, LCC*13, RMD07]. Other contributions simulate a series of low-speed touch events to enable user authentication or tangible interaction with objects placed on the screen [VG13, YCL*11]. Although the envisioned applications are very promising, current touchscreen controllers are the limiting factor. I believe that we need to rethink the capacitive touchscreen to overcome limitations in data rate, and allow for full-duplex communications and spatial information encoding. Using capacitive communication in touchscreens allows for short-range communication with physical objects through air and through the human body. Messages can be encoded locally in certain areas on the touchscreen to exploit spatial correspondences in interaction design. Possible application examples range from transferring data from displays to smartwatches or using sensor-augmented tangible objects for gaming. very interesting approaches based on air flows [SPGI13] or even electric muscle stimulation [LB13].

Figure 8.2.: Low-cost whole-body interfaces can be an important new device category in the future. In my opinion, they can be seen as a transition technology to integrated wearable interaction systems.

Towards low-cost whole-body interfaces: In my previous work I designed a variety of flat devices like Rainbowfish [GPBW14a, GPBW*14b], which can detect gestures in distances up to 30 cm. As users often experienced problems in interaction, I started to equip these devices with feedback and feed-forward mechanisms mostly based on light. I am also interested in integrating additional feedback modalities, for example based on air flows [SPGI13] or electric muscle stimulation [LB13]. In the future, I would like to make these surfaces smaller and lighter to be worn on the human body or attached to a wall for controlling a smarthome. In the simplest configuration they can be used as an on/off button or as a continuous dimmer for light. Combining multiple surfaces can be exploited for authentication with infrastructure using gestures or equipping furniture with interactive capabilities [MBRS14]. Currently, I also transfer the presented approach to the world of wearable computing in order to increase traffic safety for pedestrians, bikers and joggers. Here, I investigate the applicability of a gesture-controllable LED grid that can be attached to a user's arms, legs or body parts, as shown below. The device, shown in Figure 8.2, uses accelerometers in combination with a proximity-sensing surface to recognize the user's context for explicit and implicit interaction.

Towards an electric field identity of everyday objects, places, and devices: This idea originates from my work conducted on Capacitive Near-Field Communication [GPHW*14]. Here, smart objects emit an electric field and encode information about their manipulations. Manipulating these objects enables to naturally interact with smartwatches and other higher-level devices. Currently, I focused on *actively* sending messages from object to object, for example by transmitting information about their acceleration. However, when physical things are not equipped with communication abilities, it is still possible to *passively* identify them. This can be achieved with any device that emits or is coupled to a changing electric field [CMPT12]. Integrating this approach in smartwatches could enable to identify infrastructure or recognize devices with switch-mode power supplies. This broadens the scope of capacitive communications and provides detailed activity-related data. For example, active and passive messaging in combination with smart-metering can lead to the generation of personalized energy-consumption models [SBK*11].

Acknowledgements

I am very grateful for the support of many individuals who inspired and supported me on my way. I would like to thank Prof. Dieter W. Fellner who guided and supported me for several years. Since my time as his Bachelor's student he shaped my ways of approaching problems and sharpened my scientific senses. He often challenged my demos and caused new unintended effects and discoveries. This led me to many new ideas and opened new perspectives in thinking. I would like to thank Prof. Kristof van Laerhoven for his valuable support and his agreement to evaluate this thesis. I express my gratitude towards Prof. Arjan Kuijper who supported me on my way to this work. I benefited greatly from his broad scientific experience, his open mind, his open door, and his thorough reviews on my papers. Also, I thank him for his fantastic way of advising my Master's and Bachelor's students.

I would like to express my thanks to Reiner Wichert, who supported my scientific work even during stressful periods. He encouraged me to submit papers to various conferences and provided guidance in ideas and scientific thinking. With his broad sense for the HCI and UbiComp world, he helped me to deduct a compelling line of argument in my works. My special gratitudes go to Andreas Braun, who introduced me to capacitive sensing and made an extensive review of this dissertation. Together with Alexander Marinc, he encouraged me to write my first paper about the smart couch. This laid the foundations for my scientific curiosity and the fascination in visiting conferences. I express my thanks to Biying Fu who has been very helpful in reviewing this dissertation. I am thankful for the support of many individuals in my department, namely Tim Dutz, Felix Kamieth, Cornelia Kurkowski, Martin Majewski, Carsten Stocklöw, and Saied Tazari.

In my work, I cooperated closely with researchers and practitioners from other institutions. I would like to thank Eugen Berlin, Marko Borazio, and Frank Englert from TU Darmstadt. Raphael Wimmer from University of Regensburg has been an extraordinary source of discussion and excursion. I thank JW Nickel from RocketScience San Francisco for his help on including results from Eric Staller's art installation. Hansjörg Lienert from Dräger & Lienert helped me a lot to understand the requirements on tangible interaction for visually impaired people. I would like to thank Prof. Patrick Baudisch, Jingyuan Cheng, Euan Freeman, Prof. Marco Gruteser, Christian Holz, Pranav Mistry, Stuart Taylor, and Raphael Wimmer for permission to use figures in this thesis. Moreover, I gratefully acknowledge the effort of many anonymous researchers, who invested a lot of time in reviewing and thus improving my publications.

I had the pleasure to supervise many inspiring students with their Bachelor's and Master's thesis - namely Oskar Bechtold, Sebastian Beck, Yannick Berghöfer, Xavier Dellangnol, Pattrick Gottschämmer, Sebastian Herber, Florian Kirchbuchner, Alexander Pavlov, Lukas Strassel, Silvia Rus, Julian von Wilmsdorff, and Steeven Zeiss. I am very grateful for the opportunity of working with many extraordinary students and HiWis: Sandra Ade, Sebastian Benchea, Daniel Borrman, Nicolai Engelmann, Thorben Hehmann, David Hilkert, Simon Reinhardt, Sönke Schmidt, Christian Schuster, Gee Fung Sit, Stefan Wegener, Pia Weiss, Justus Weiss, and Daniel Wilbers, and Sinaly Wonogo. Among with many other persons mentioned above, I studied Information System Technology at TU Darmstadt. My gratitude goes to Prof. Andy Schürr for supporting us all in such constructive manner and always having an open ear to the students' belongings and wishes.

My family has been especially helpful and encouraged me on my way. My parents always supported me during my studies and had the faith that I will find my way. I would like to thank Korinna Allhoff for her support and patience, especially in stressful times. She managed to pull me away from my computer to enjoy the sunny sides of social life. Special thanks go to Monika Jasper-Allhoff who reviewed many papers and improved my English skills.

A. Hardware Schematics

A.1. Honeyfish

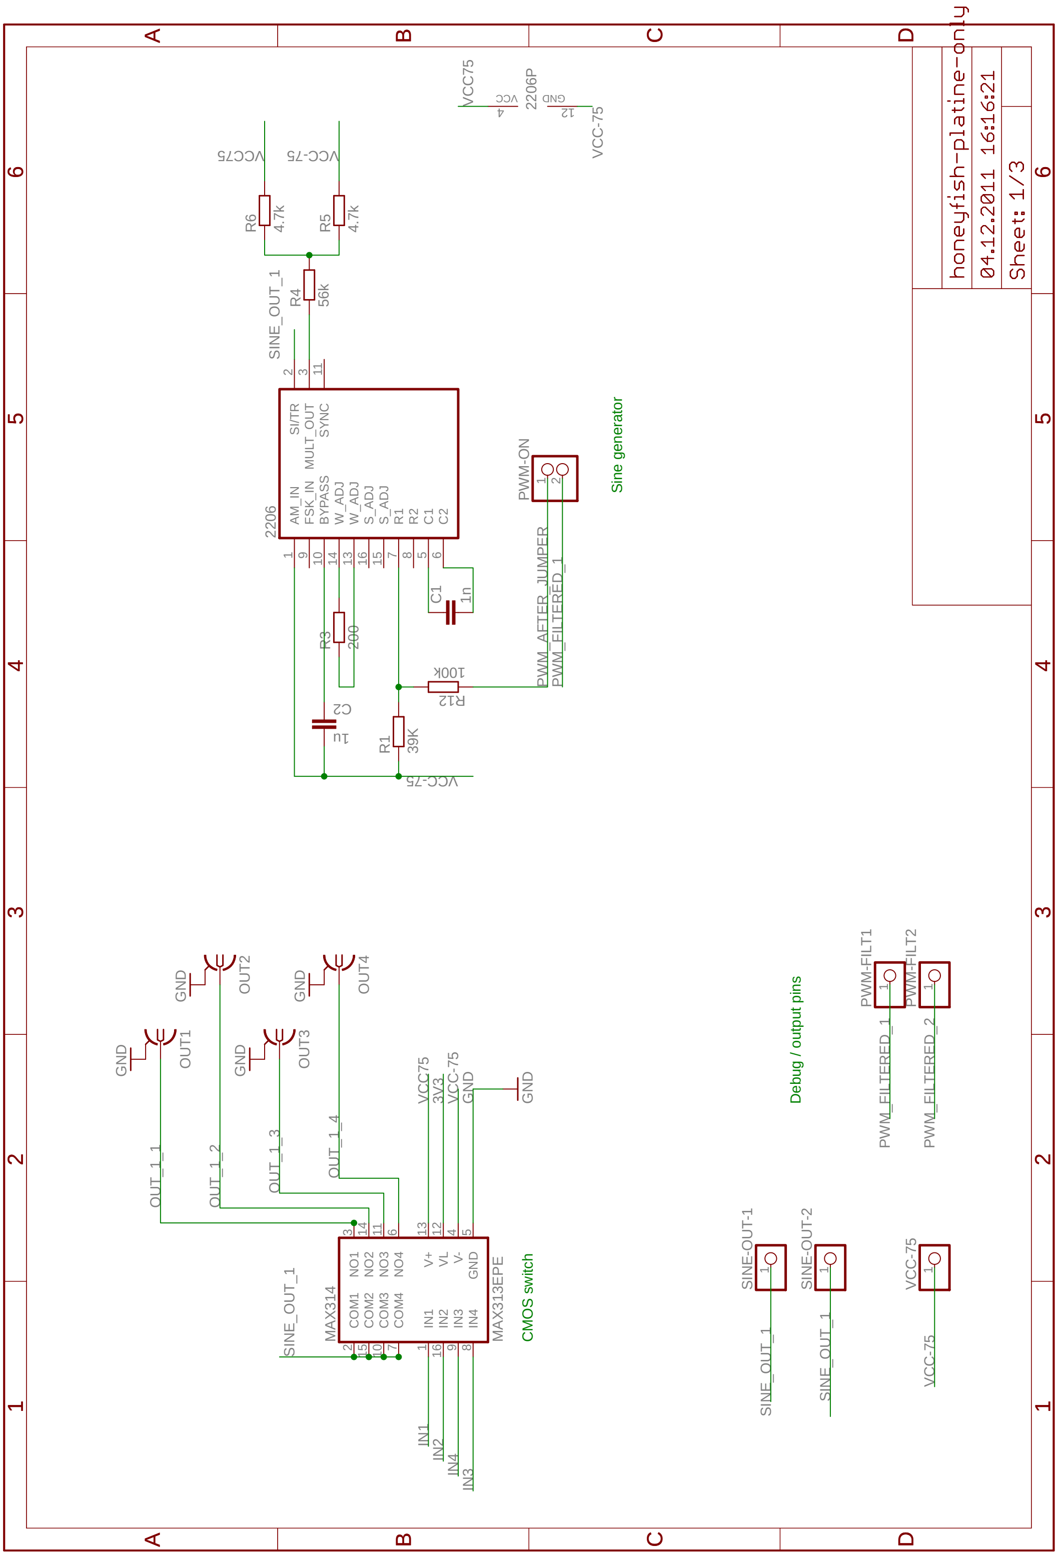

Figure A.1.: Honeyfish board schematic (page 1).

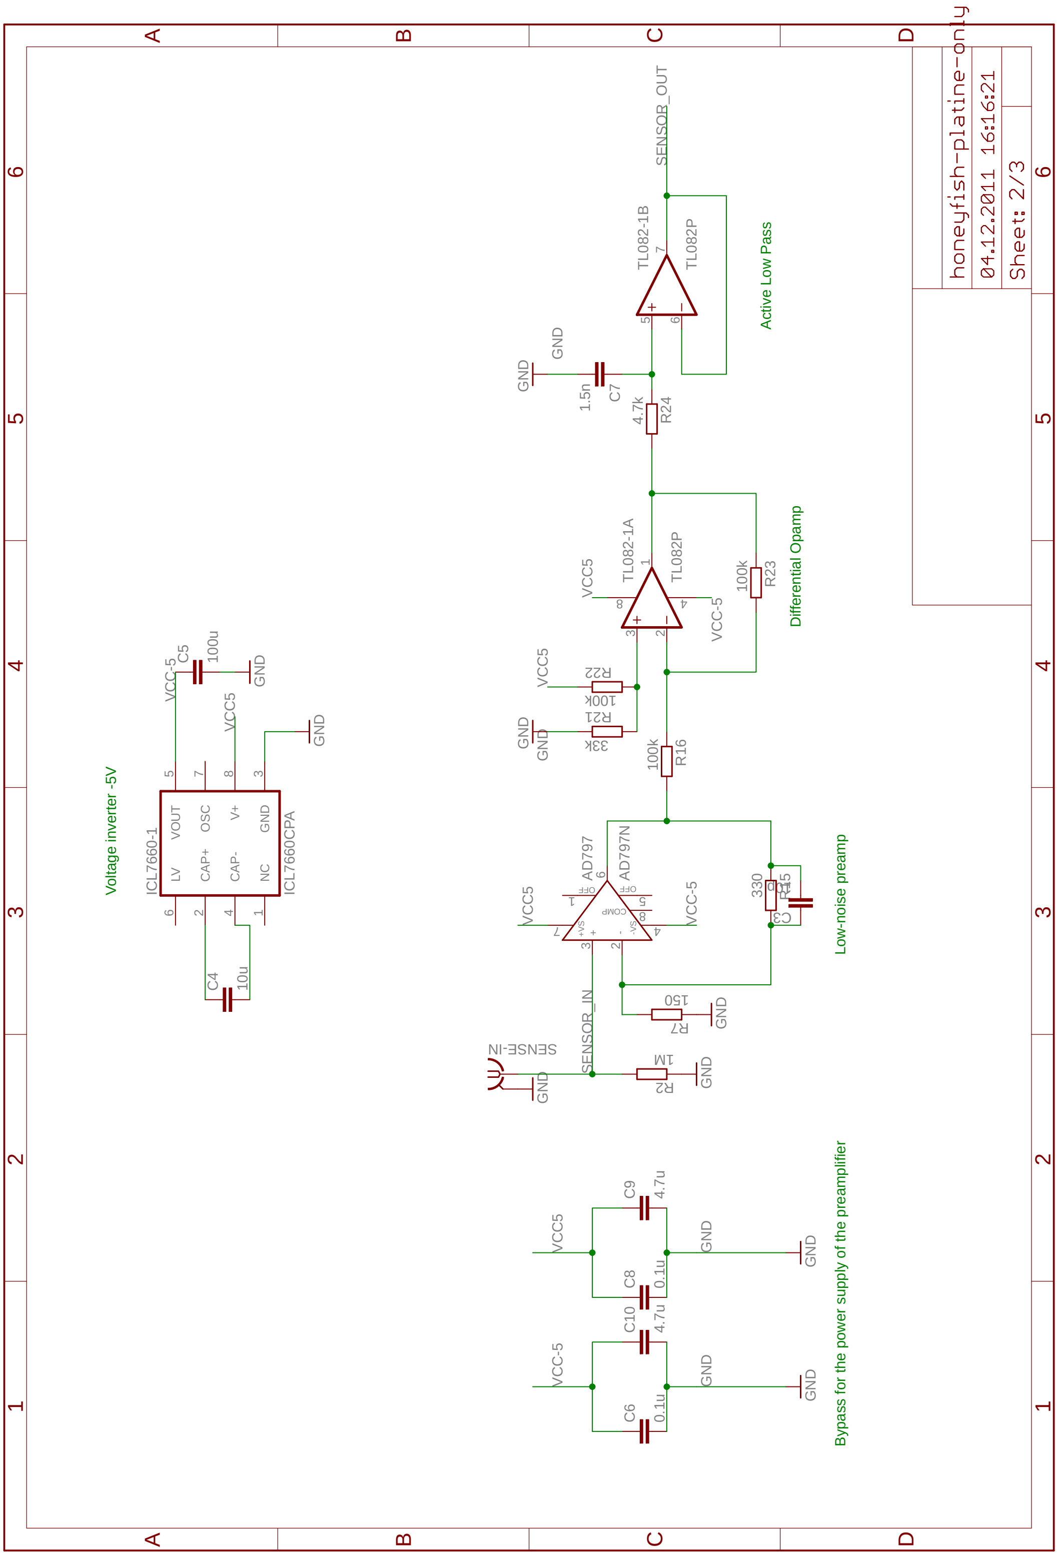

Figure A.2.: Honeyfish board schematic (page 2).

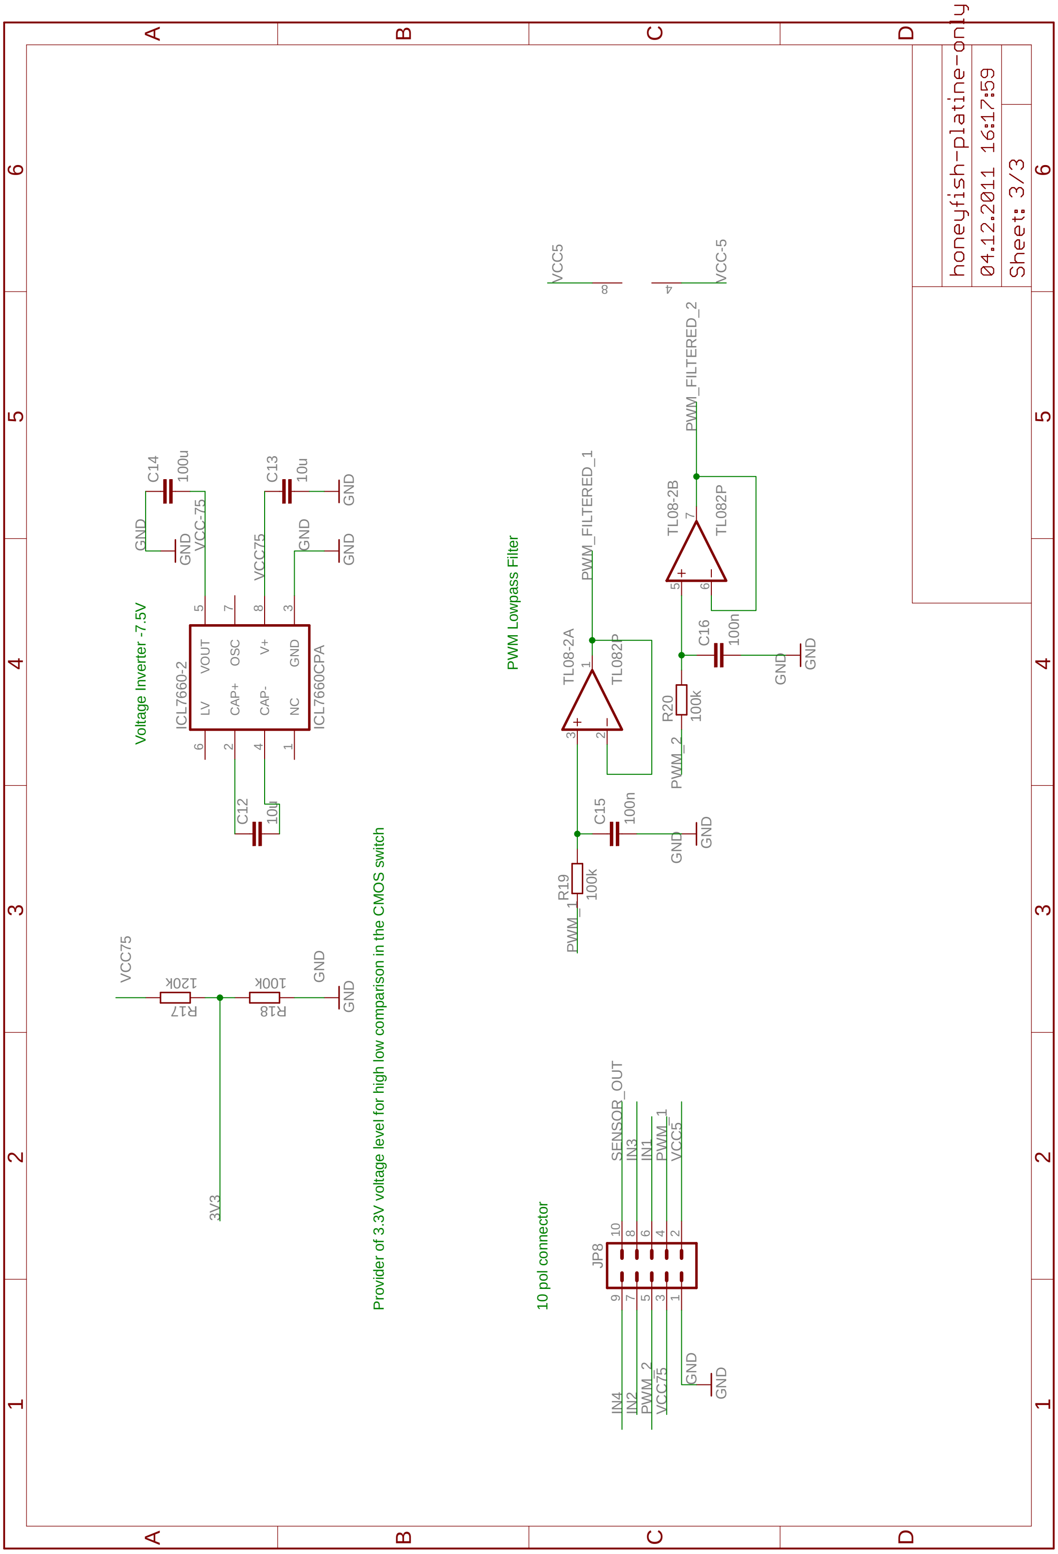

Figure A.3.: Honeyfish board schematic (page 3).

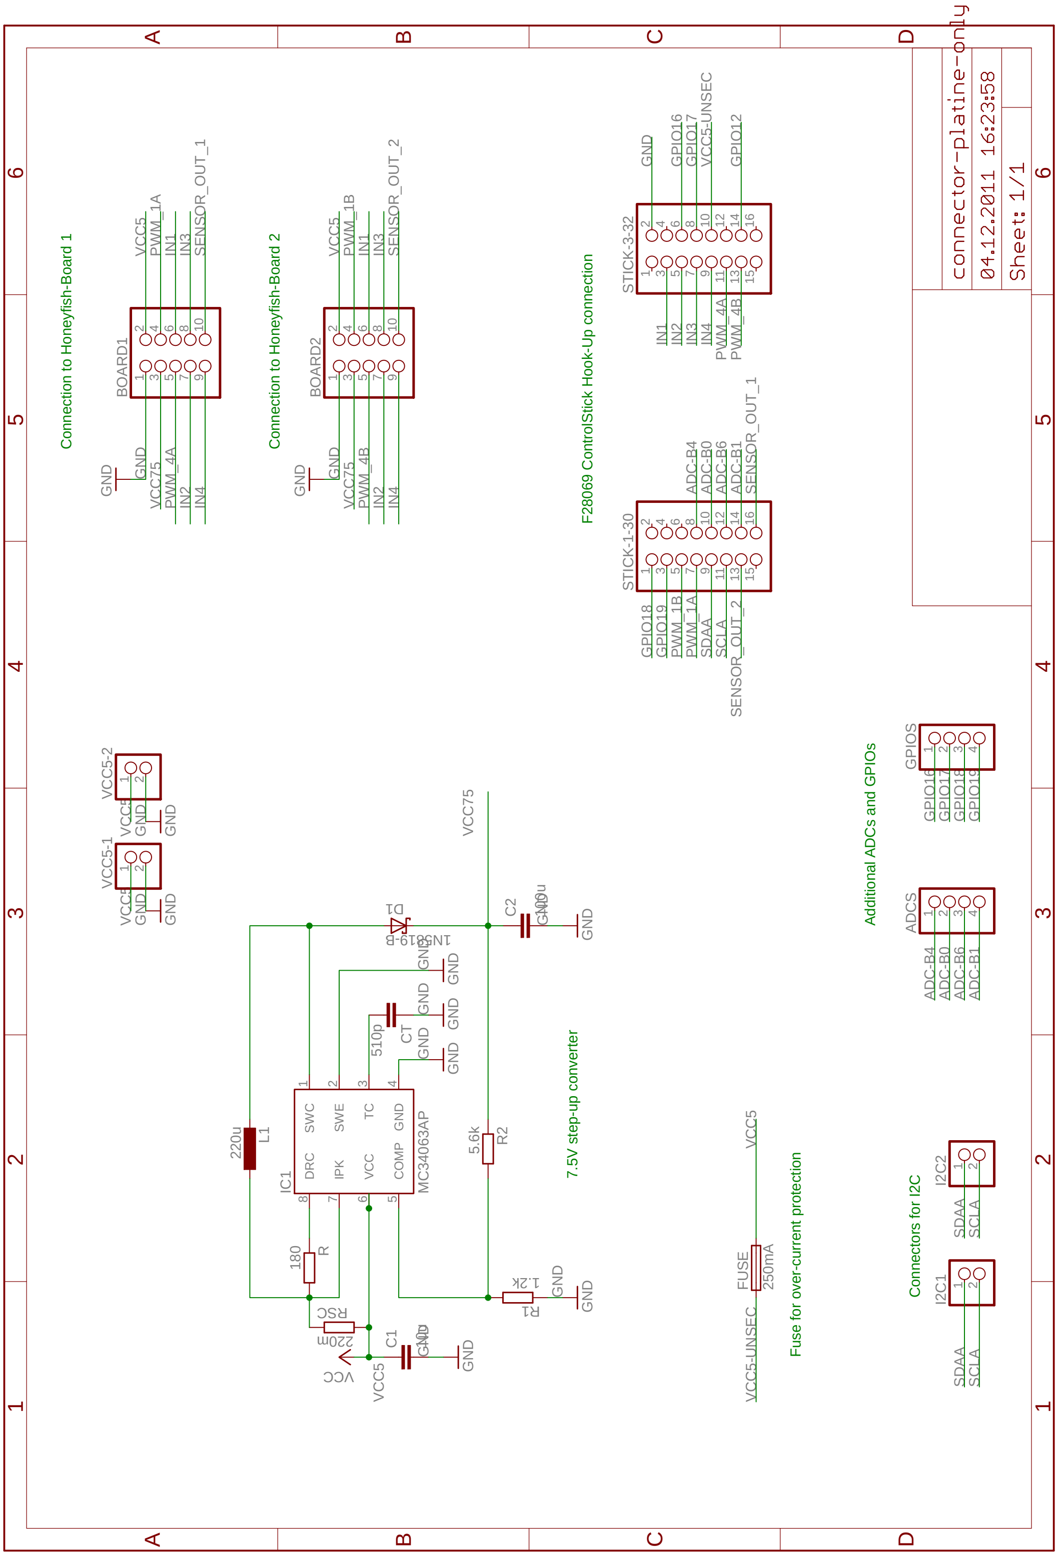

Figure A.4.: Connector board schematic (page 1).

A.2. Rainbowfish

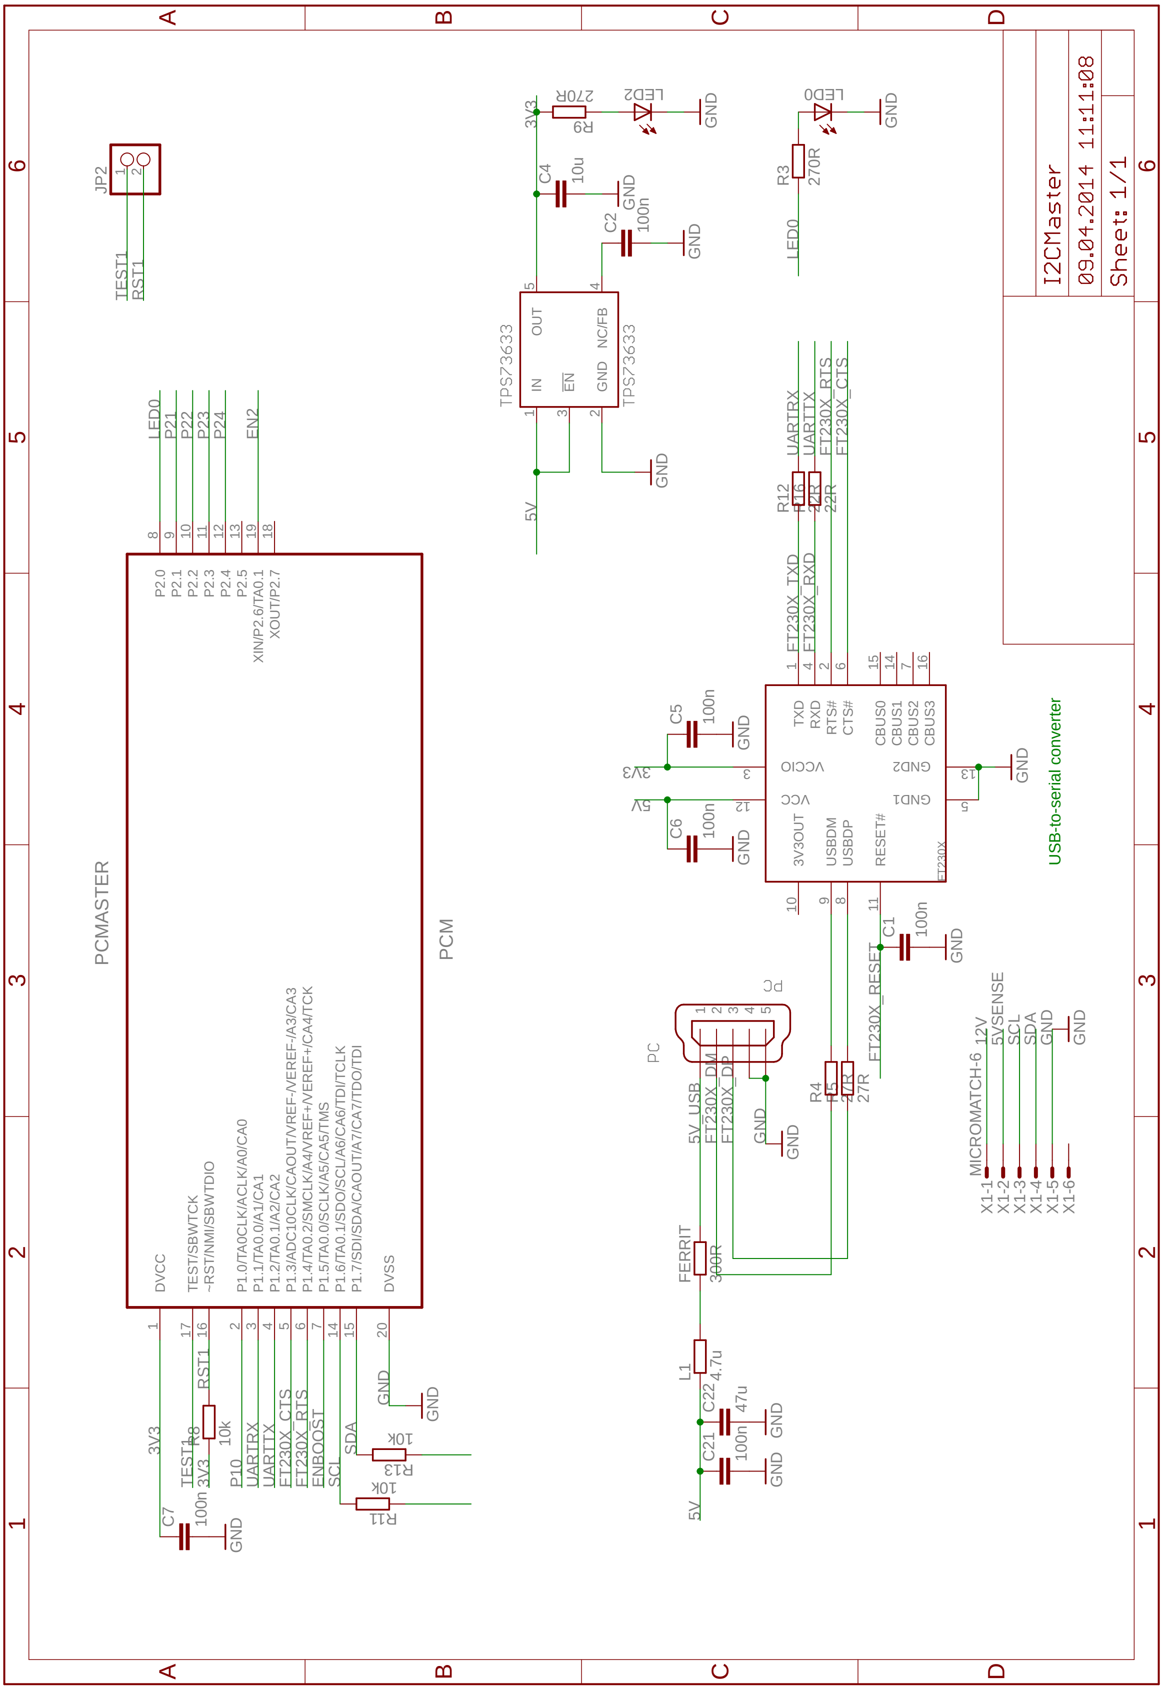

Figure A.5.: Rainbowfish master board schematic (page 1).

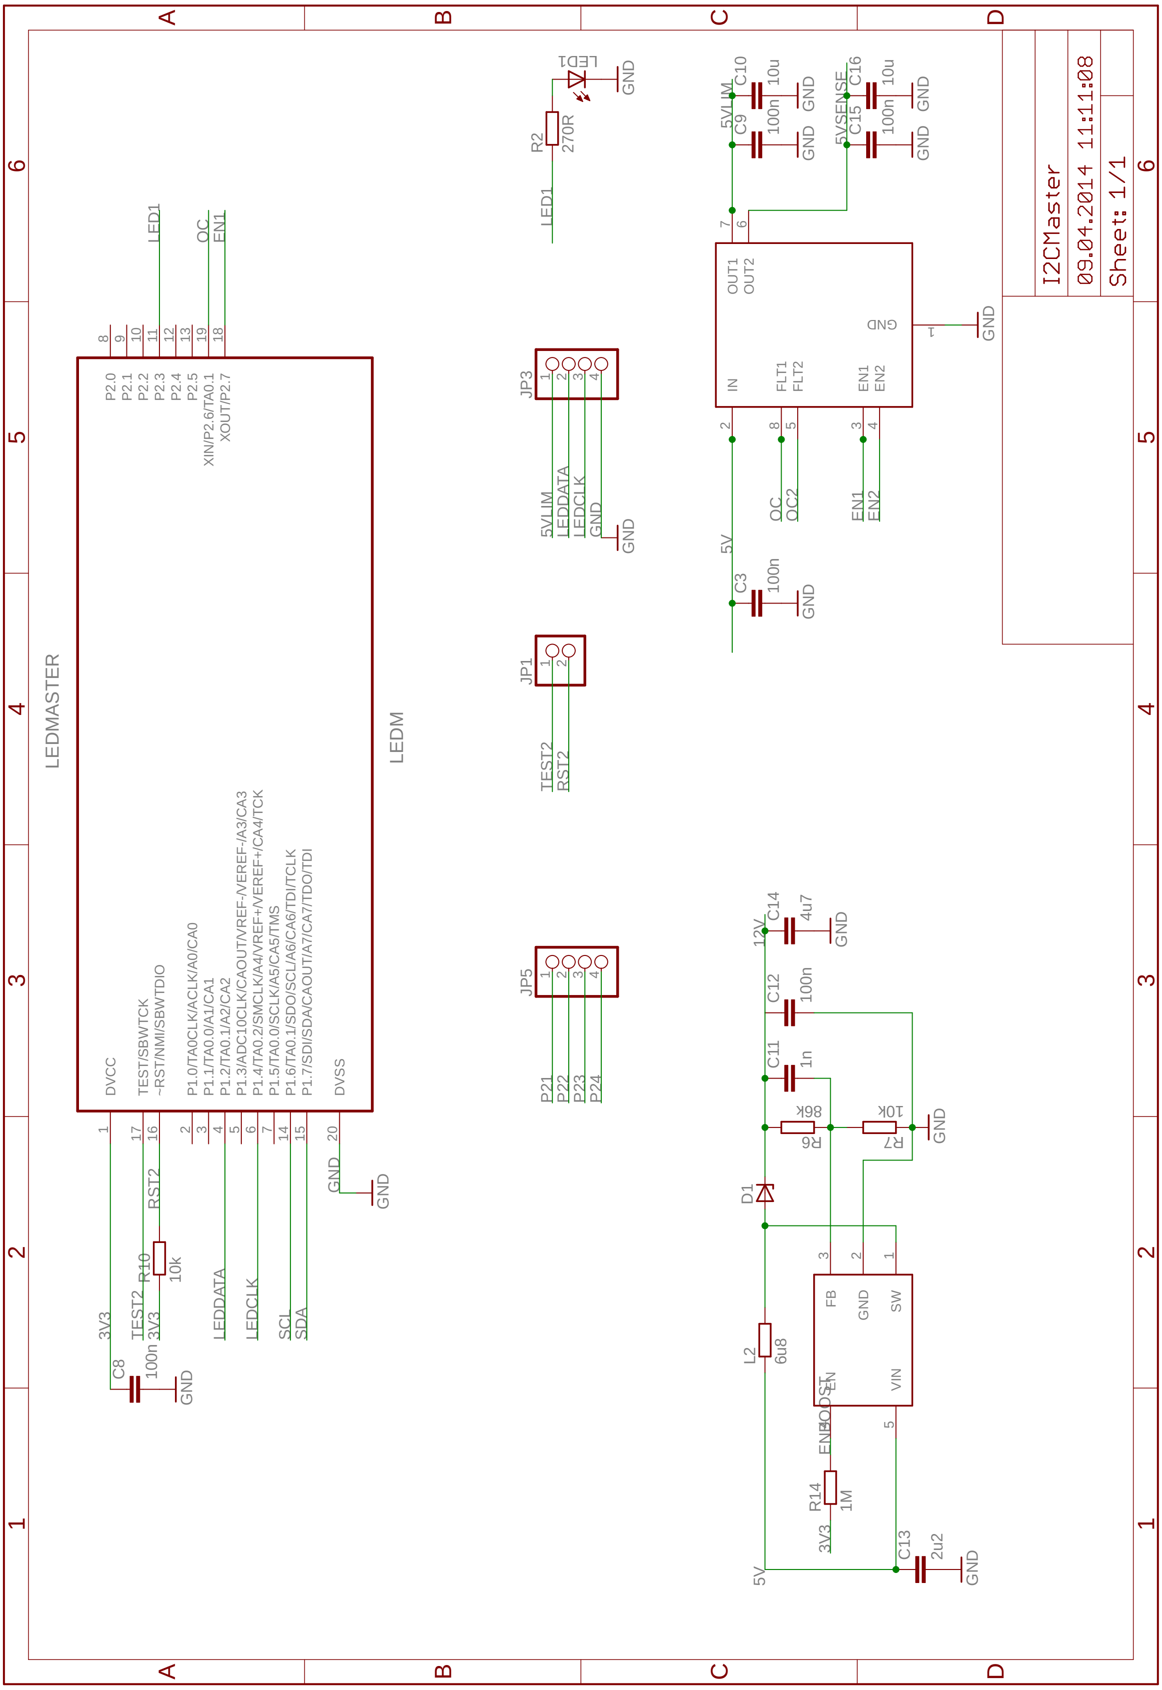

Figure A.6.: Rainbowfish master board schematic (page 2).

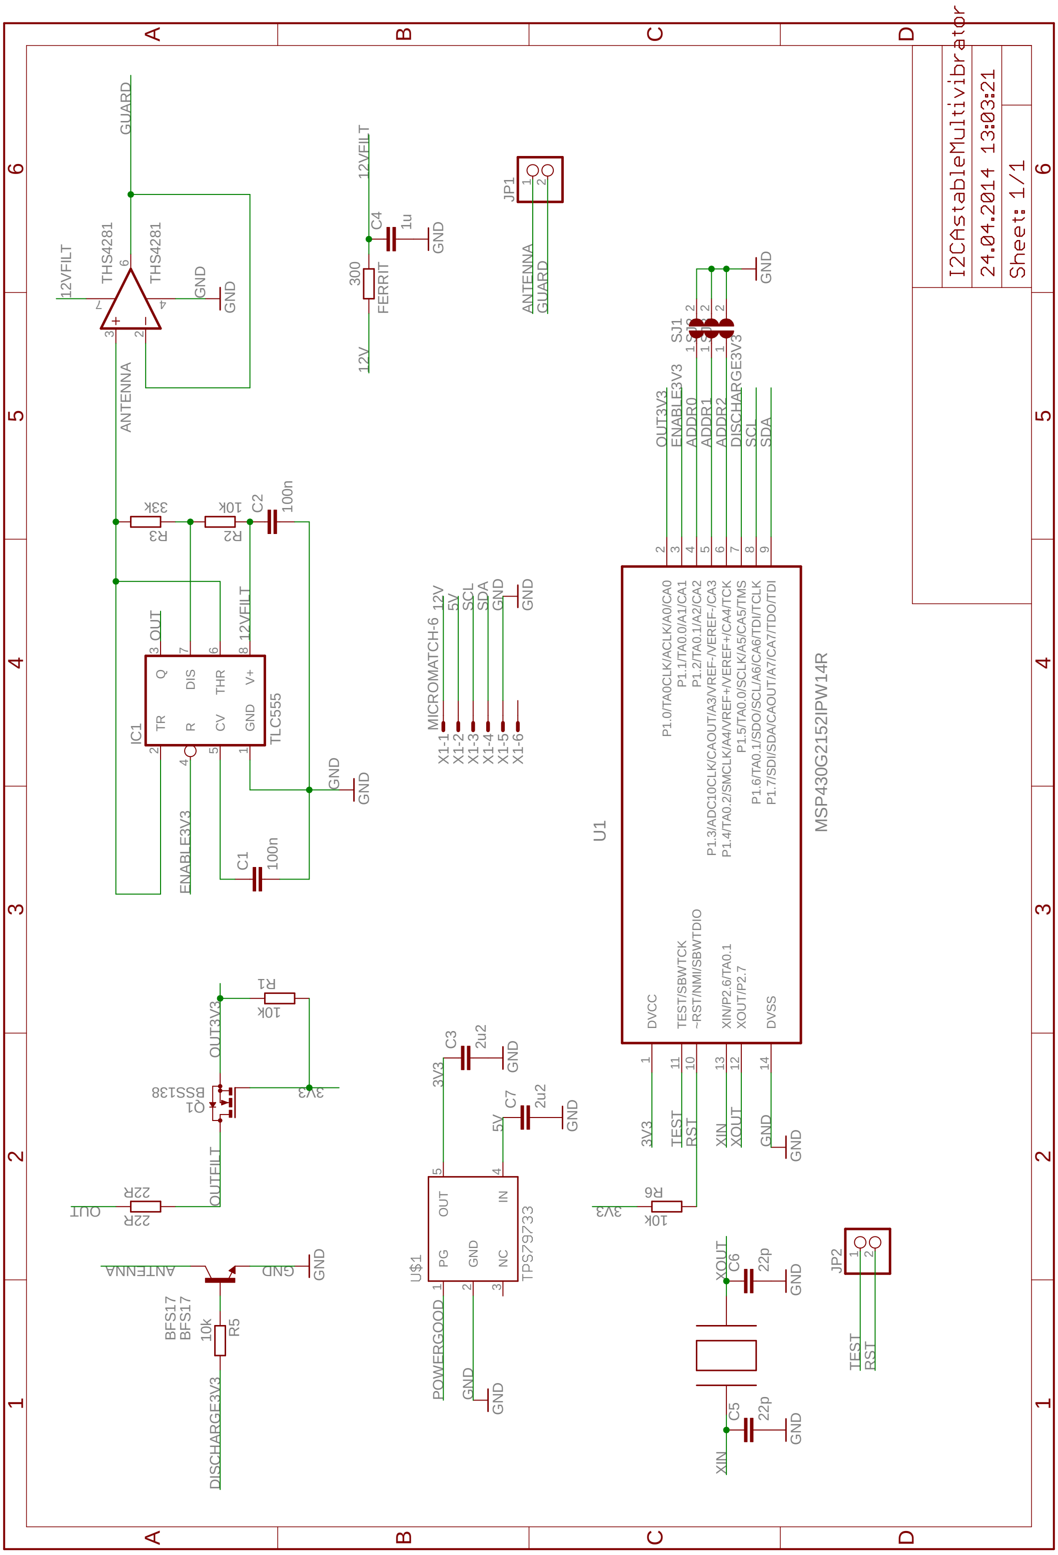

Figure A.7.: Rainbowfish sensor board schematic (page 1).

A.3. OpenCapSense

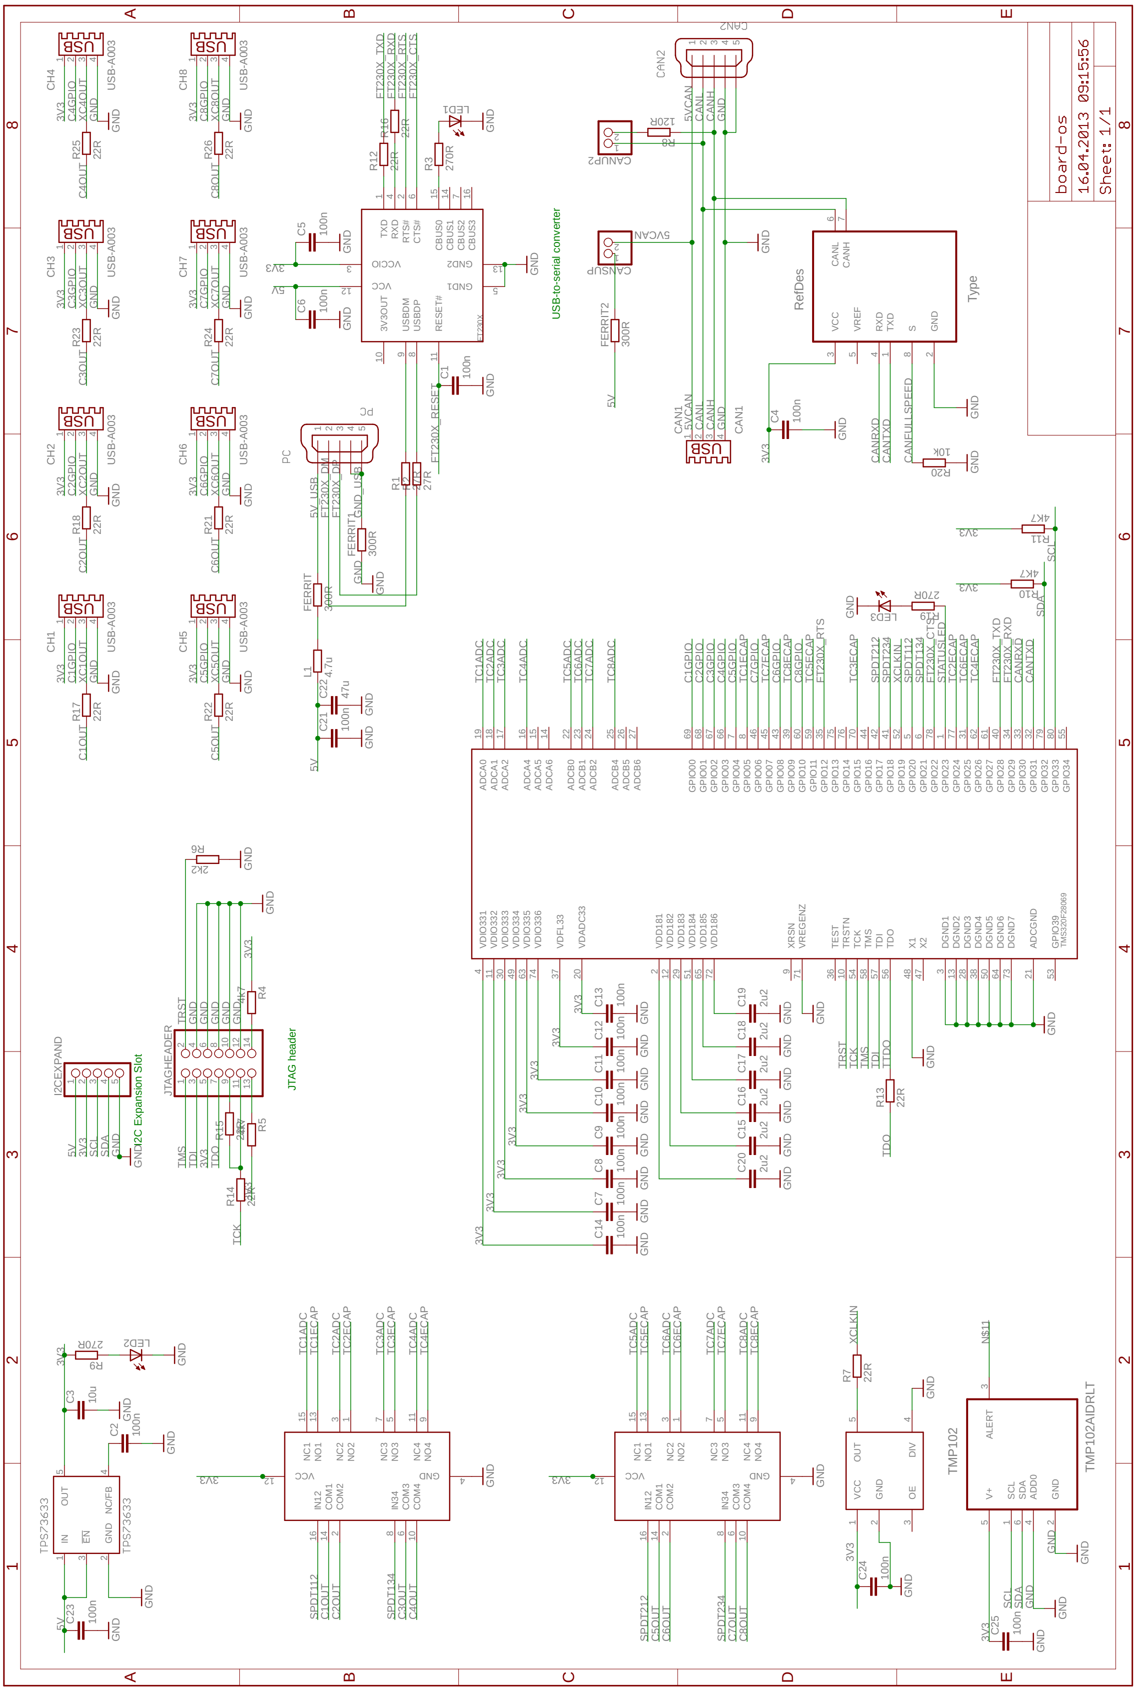

Figure A.8.: OpenCapSense controller board schematic.

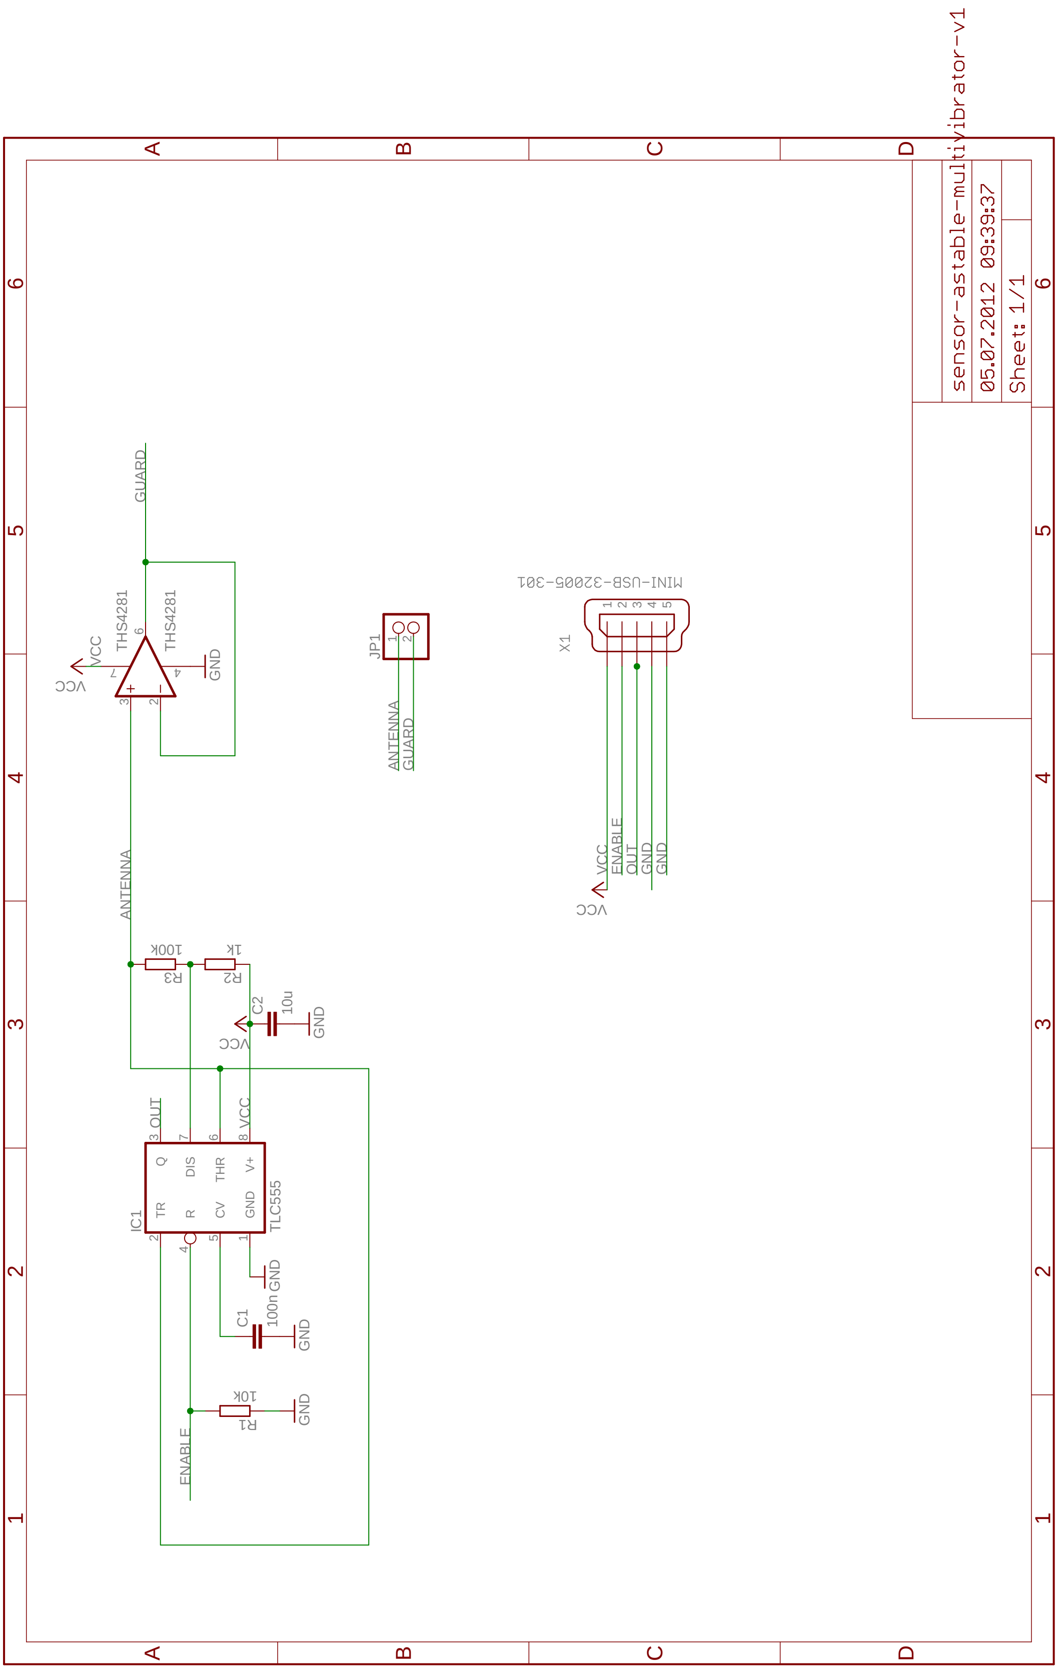

Figure A.9.: OpenCapSense loading mode sensor.

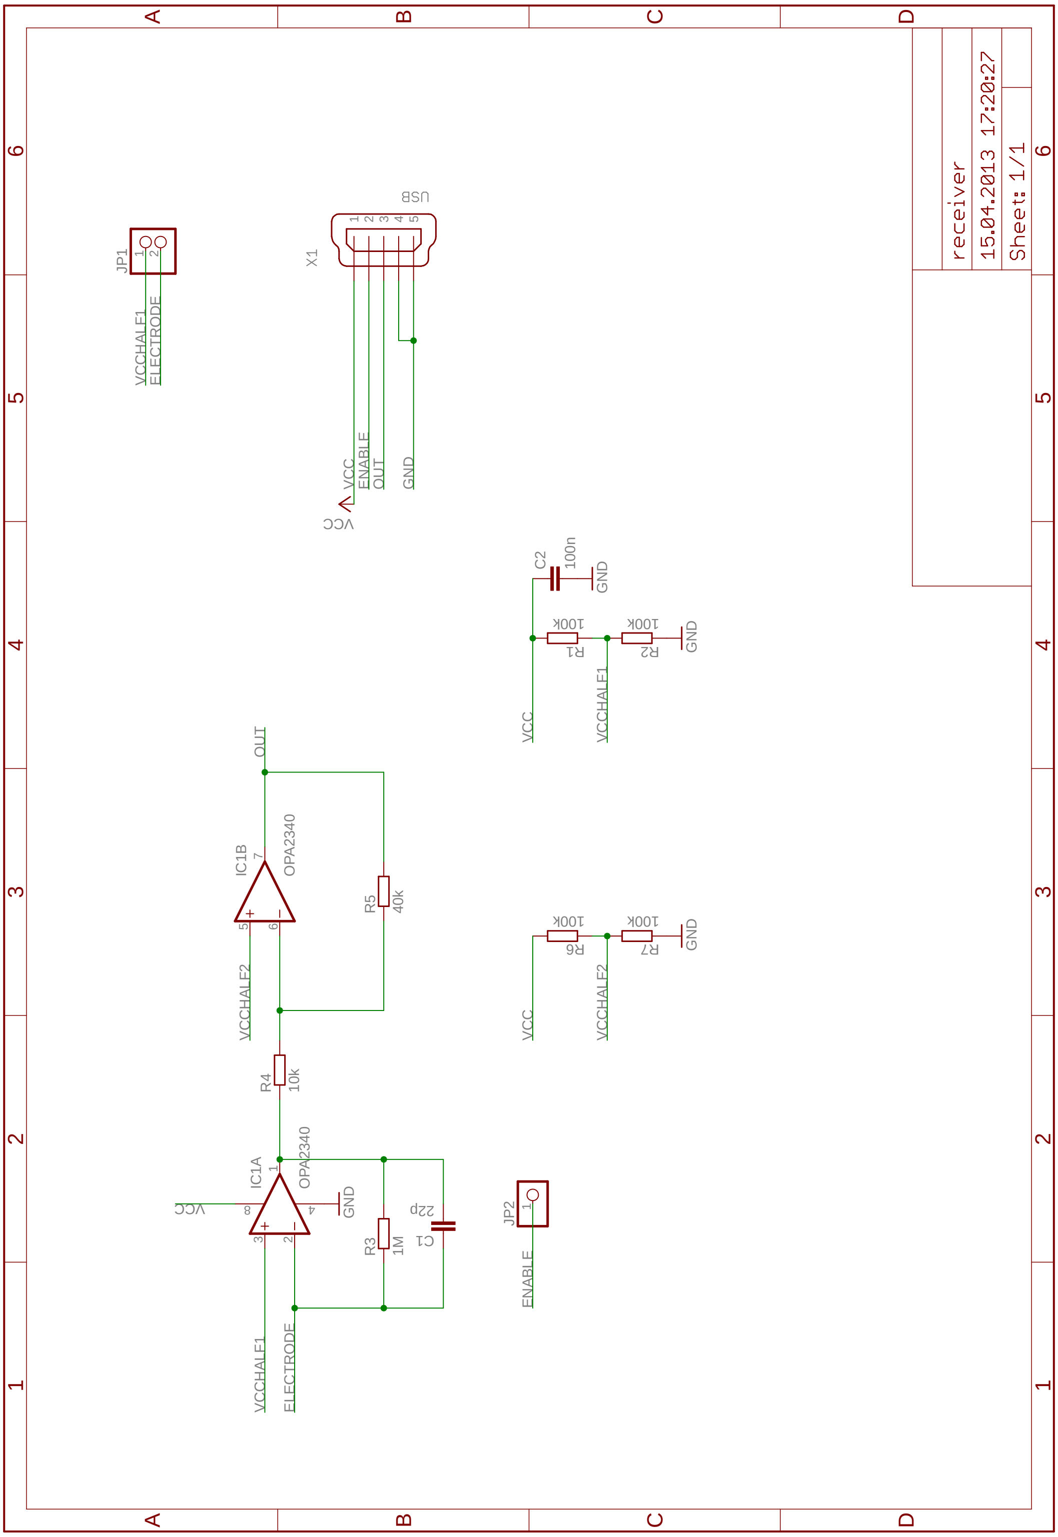

Figure A.10.: OpenCapSense shunt mode sensor.

A.4. CapNFC

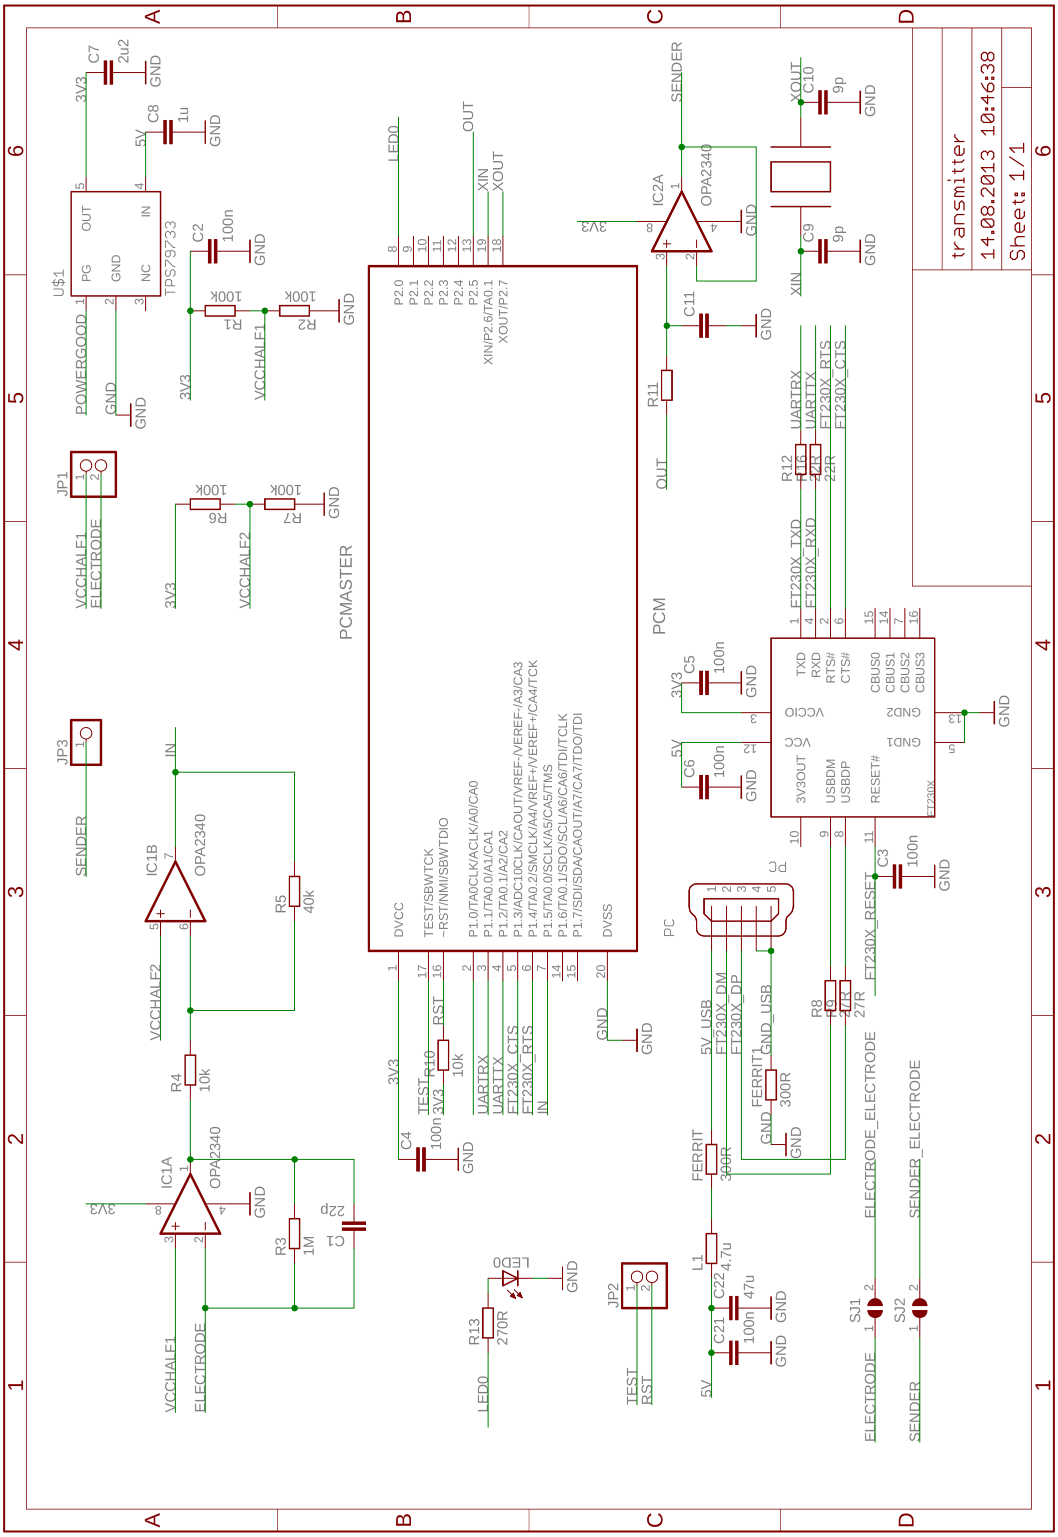

Figure A.11.: CapNFC transceiver board schematic.

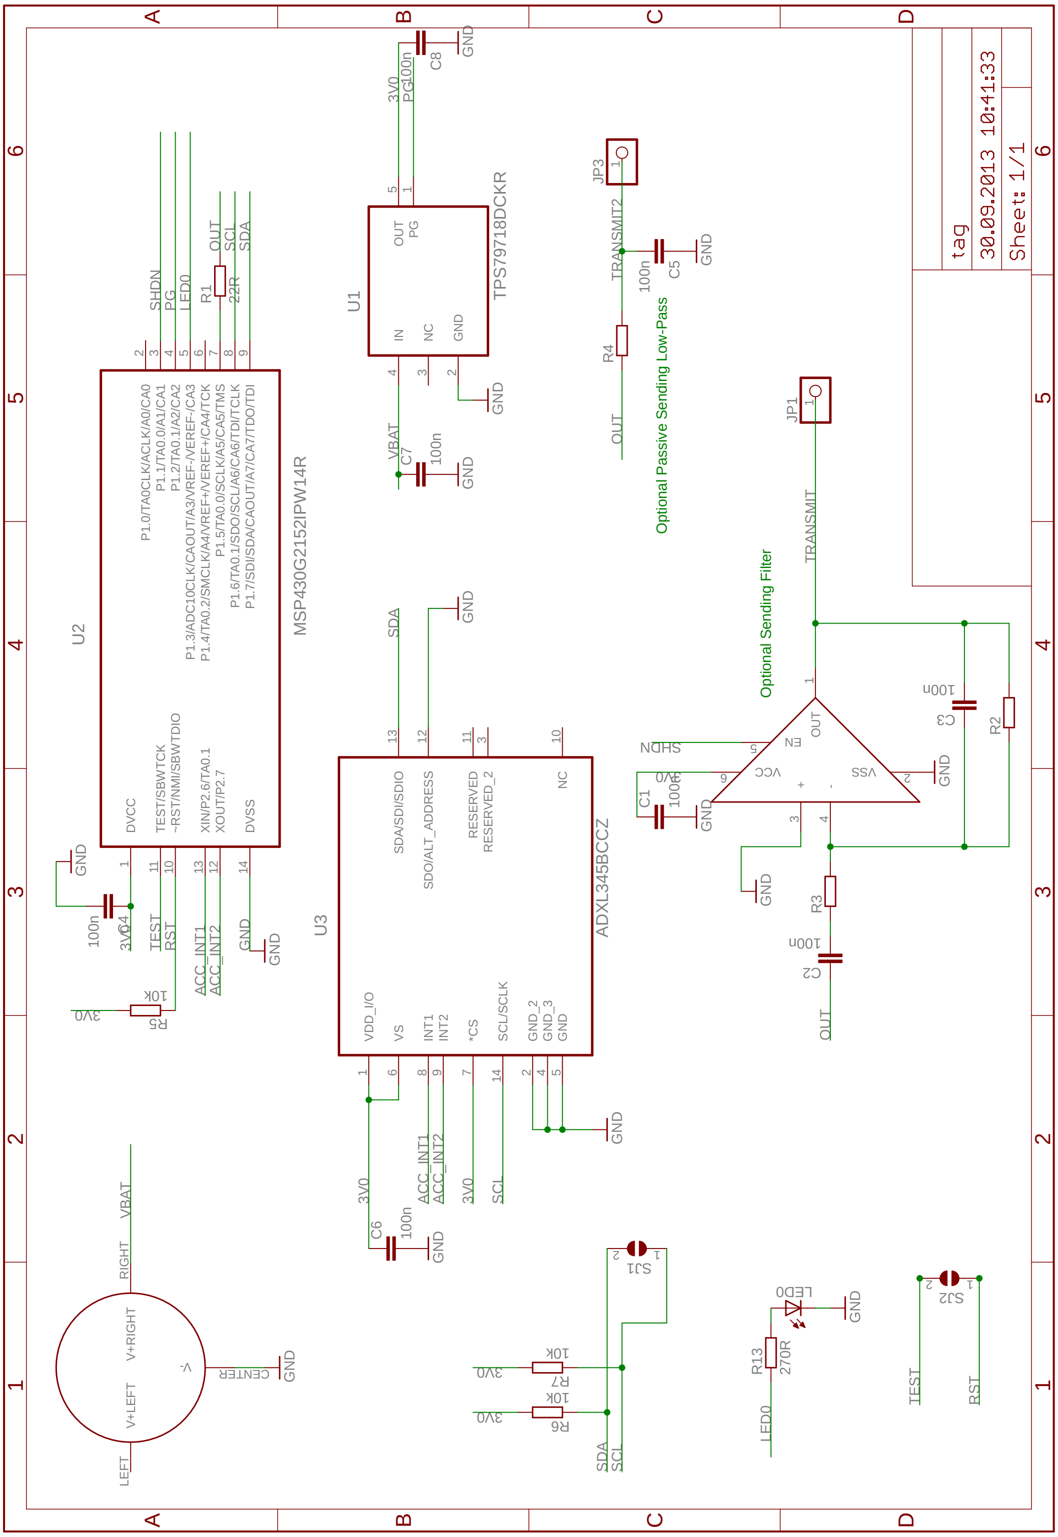

Figure A.12.: CapNFC tag with accelerometer.

B. Publications and Talks

The thesis is partially based on the following publications and talks:

Full Conference Papers

C.12 Gottschämmer, S., **Grosse-Puppendahl, T.**, Kuijper, A.: User Location Modeling based on Heterogeneous Data Sources. In: Distributed, Ambient, and Pervasive Interactions 2015. Lecture Notes in Computer Science. Springer (to appear) (2015)

C.11 Fu, B., **Grosse-Puppendahl, T.**, Kuijper, A.: A Gesture Recognition Method for Proximity-Sensing Surfaces in Smart Environments. In: Distributed, Ambient, and Pervasive Interactions 2015. Lecture Notes in Computer Science. Springer (to appear) (2015)

C.10 Rus, S., **Grosse-Puppendahl, T.**, Kuijper, A.: Recognition of Bed Postures using Mutual Capacitance Sensing. In: Ambient Intelligence. Lecture Notes in Computer Science. Springer (2014)
→ Best Paper Award

C.9 **Grosse-Puppendahl, T.**, Herber, S., Wimmer, R., Englert, F., Beck, S., Wichert, R., Kuijper, A.: Capacitive Near-Field Communication for Ubiquitous Interaction and Perception. In: 2014 ACM International Joint Conference on Pervasive and Ubiquitous Computing – UbiComp. ACM (2014)
→ Acceptance Rate: 20%, Best Paper Nominee

C.8 **Grosse-Puppendahl, T.**, Beck, S., Wilbers, D., Zeiss, S., von Wilmsdorff, J., Kuijper, A.: Ambient Gesture-Recognizing Surfaces with Visual Feedback. In: Distributed, Ambient, and Pervasive Interactions 2014. Lecture Notes in Computer Science. Springer (2014)

C.7 Zeiss, S., Marinc, A., Braun, A., **Grosse-Puppendahl, T.**, Beck, S.: A Gesture-based Door Control using Capacitive Sensors. In: Distributed, Ambient, and Pervasive Interactions 2014. Lecture Notes in Computer Science. Springer (2014)

C.6 **Grosse-Puppendahl, T.**, Benchea, S., Kamieth, F., Braun, A., Schuster, C.: Unobtrusive Recognition of Working Situations. In: Distributed, Ambient, and Pervasive Interactions 2013. Lecture Notes in Computer Science. Springer (2013)

C.5 **Grosse-Puppendahl, T.**, Braun, A., Kamieth, F., Kuijper, A.: Swiss-cheese Extended: An Object Recognition Method for Ubiquitous Interfaces based on Capacitive Proximity Sensing. In: Proceedings of the SIGCHI Conference on Human Factors in Computing Systems. ACM (2013)
→ Acceptance Rate: 20%

C.4 **Grosse-Puppendahl, T.**, Berghoefer, Y., Braun, A., Wimmer, R., Kuijper, A.: OpenCapSense: A Rapid Prototyping Toolkit for Pervasive Interaction using Capacitive Sensing. In: IEEE International Conference on Pervasive Computing and Communications. IEEE (2013)
→ Acceptance Rate: 15%

C.3 **Grosse-Puppendahl, T.**, Berlin, E., Borazio, M.: Enhancing Accelerometer-based Activity Recognition with Capacitive Proximity Sensing. In: Ambient Intelligence. Lecture Notes in Computer Science, pp. 17-32. Springer (2012)

C.2 **Grosse-Puppendahl, T.**, Braun, A.: Honeyfish - A High Resolution Gesture Recognition System based on Capacitive Proximity Sensing. In: EmbeddedWorld Conference. Weka Fachmedien (2012)

C.1 **Grosse-Puppendahl, T.**, Marinc, A., Braun, A.: Classification of User Postures with Capacitive Proximity Sensors in AAL-Environments. In: Ambient Intelligence. Lecture Notes in Computer Science, vol. 7040, pp. 314-323. Springer (2011)

Workshop Papers & Demonstrations

W.2 Fu, B., Karolus, J., **Grosse-Puppendahl, T.**, Hermann, J., Kuijper, A.: Opportunities for Activity Recognition using Ultrasound Doppler Sensing on Unmodified Mobile Phones. In: iWOAR 2015 – 2nd international Workshop on Sensor-based Activity Recognition and Interaction. ACM (to appear) (2015)

W.1 **Grosse-Puppendahl, T.**, Weiss, J., Weiss, P., Herber, S., Lienert, H.: Smart Objects in Accessible Warehouses for the Visually Impaired. In: Third Workshop on Smart Objects in conjunction with ACM 2014 International Conference on Intelligent User Interfaces. TU Prints (2014)

DM.2 Liang, R. and Chan, L. and Tseng, H. and Kuo, H. and Huang, D. and Yang, D. and Chen, B. and **Grosse-Puppendahl, T.** and Beck, S. and Wilbers, D. and Kuijper, A. and Heo, H. and Park, H. and Kim, S. and Chung, J. and Lee, G. and Lee, W. and Unander-Scharin, C. and Unander-Scharin, A. and Hook, K. and Elblaus, L.: Demo Hour. In: ACM interactions magazine. 21 (5), 6-9. ACM (2014)

DM.1 **Grosse-Puppendahl, T.**, Beck, S., Wilbers, D.: Rainbowfish: Visual Feedback on Gesture-Recognizing Surfaces. In: CHI '14 Extended Abstracts on Human Factors in Computing Systems. ACM (2014)

Other Contributions

NO.1 Banhatti, R. D., Brylok, A., Doser, M., Dreiner, T., **Grosse-Puppendahl, T.**, Hoppe, A., Joska, R., Laurila-Dürsch, J., Lauterbach, C., Ludwig, T., Reiß, C., Schaper, A., Schirp, C., Schliepkorte, H., Tiedtke, S.: Technikunterstütztes Leben - Ambient Assisted Living (AAL) - Prozessunterstützung zur technischen Realisierung von Assistenzsystemen (umgebungsunterstützender Technik) in Gebäude und Wohnumfeld (in German). VDE Verlag (2014)

Working Papers

WP.1 **Grosse-Puppendahl, T.**, Bechtold, O., Strassel, L., Kuijper, A.: Enhancing Traffic Safety with Wearable Low-Resolution Displays

WP.2 Kirchbuchner, F., **Grosse-Puppendahl, T.**, Hastall, M., Distler, M., Kuijper, A.: Smart Living Innovations from Senior Citizens' Perspectives: Understanding Privacy Concerns, Technology Acceptance, and Expectations

Theses

D.1 **Grosse-Puppendahl, T.**: Capacitive Sensing and Communication for Ubiquitous Interaction and Perception, Ph.D. thesis, Technische Universität Darmstadt, Germany (04/2015)

DT.2 **Grosse-Puppendahl, T.**: Multi-hand Interaction using Custom Capacitive Proximity Sensors, Master's thesis, Technische Universität Darmstadt, Germany (04/2012)

DT.1 **Grosse-Puppendahl, T.**: Sensor-based Activity Visualization for Monitoring Daily Schedules, Bachelor's thesis, Technische Universität Darmstadt, Germany (04/2010)

Invited Talks

IT.7 Sensors and Devices Group, Microsoft Research, Cambridge, UK
Invited talk: Supporting Proxemic Interactions with Multi-Scale Electric-Field Sensing, hosted by Steve Hodges (03/2015)

IT.6 Human-Computer Interaction Group, TELECOM ParisTech, Paris, France
Invited talk: Capacitive Sensing and Communciation for Ubiquitous Interaction and Perception, hosted by Gilles Bailly (02/2015)

IT.5 Sensor Technology Research Center, University of Sussex, Brighton, UK
Invited talk: Capacitive Sensing and Communciation for Ubiquitous Interaction and Perception, hosted by Daniel Roggen and Robert Prance (11/2014)

IT.4 WinLab, Rutgers University, Newark, USA
Invited talk: Capacitive User Interfaces in Smart Environments, hosted by Marco Gruteser (09/2014)

IT.3 Technische Universität Dortmund, Dortmund, Germany
Faculty of Rehabilitation Sciences, Language and Communication
Invited Lecture: Technische Assistenzsysteme für ältere Menschen und Menschen mit Behinderungen (in German), hosted by Matthias R. Hastall (05/2014)

IT.2 Fraunhofer IDM@NTU, Nanyang Technical University, Singapore
Invited Talk: Ubiquitous Interaction and Perception using Capacitive Sensing, hosted by Wolfgang Müller-Wittig (01/2014)

IT.1 Technische Universität Dortmund, Dortmund, Germany
Faculty of Rehabilitation Sciences, Language and Communication
Invited Talk: Communication technologies for Persons suffering from Neurodegenerative Diseases - Current Developments in Ambient-Assisted-Living Research, hosted by Matthias R. Hastall (11/2012)

C. Supervising Activities

The following list summarizes the student bachelor, diploma and master thesis supervised by the author. The results of these works were partially used as an input into the thesis.

Supervision & Teaching

Master's theses

03/2015 - 09/2015	Xavier Dellangnol (Technische Universität Darmstadt, Electrical Engineering): Indoor Localization of Humans based on Electric Potential Sensing
11/2014 - 04/2015	Julian von Wilmsdorff (Technische Universität Darmstadt, Information System Technology): Electric Potential Sensing in Ubiquitous Computing
09/2014 - 03/2015	Alexander Pavlov (Technische Universität Darmstadt, Electrical Engineering): Ubiquitous Activity Recognition using Wireless Capacitive Sensing Nodes
04/2014 - 10/2014	Florian Kirchbuchner (Technische Universität Darmstadt, Computer Science): User Tracking and Behavior Analysis based on a Capacitive Indoor Localization System
03/2013 - 10/2013	Silvia Rus (Technische Universität Darmstadt, Electrical Engineering): Recognition of Lying Postures using Capacitive Proximity Sensing → Fraunhofer IGD Best Thesis Award
07/2012 - 01/2013	Yannick Berghoefer (Technische Universität Darmstadt, Information System Technology): Human-Machine-Interfaces in Automotive Environments using Capacitive Proximity Sensors

Bachelor's theses

09/2014 - 03/2015	Lukas Strassel (Technische Universität Darmstadt, Computer Science): Sensing Interactions on Body-worn Low-resolution LED-Displays
09/2014 - 03/2015	Oskar Bechtold (Technische Universität Darmstadt, Computer Science): An App-driven Interaction Concept for Body-worn Low-resolution LED-Displays
01/2014 - 06/2014	Patrick Gottschämmer (Technische Universität Darmstadt, Computer Science): User Location Modelling based on Heterogeneous Data Sources
08/2013 - 11/2013	Sebastian Herber (Technische Universität Darmstadt, Information System Technology): Tangible Interaction using Capacitive Near-Field Communication
05/2013 - 07/2013	Sebastian Beck (Technische Universität Darmstadt, Information System Technology): A Gesture Recognition Device with Visual Feedback using Capacitive Proximity Sensing → Fraunhofer IGD Best Thesis Award
08/2012 - 11/2012	Steeven Zeiss (Technische Universität Darmstadt, Electrical Engineering): Development of a Contactless Closing Mechanism for Automatic Doors

D. Curriculum Vitae

Tobias Grosse-Puppendahl
Kiesstrasse 24
64283 Darmstadt
Germany

Contact Information

Telephone:	+49 6151 155565
Mobile:	+49 171 6506806
E-Mail:	tobias(at)grosse-puppendahl(dot)com
Date of Birth:	6th April 1986
Place of Birth:	Muenster
Nationality:	German
Languages:	German: native English: fluent in spoken and written French: good Dutch: basic

Work Experience

Since 04/2012 — **Fraunhofer Institute for Computer Graphics Research IGD**
Researcher and PhD candidate

- Technical project leader: Development of a smart floor for people localization and emergency recognition (e.g. falls by the elderly)
- Technical project leader: Development of an unobtrusive capacitive wireless health sensor
- Research projects on the development of embedded sensing systems, including activity and behavior recognition, gesture recognition, whole-body interaction
- Supervision of students in lectures, practical exercises and thesis

05/2004 - 12/2012 — **Grosse-Puppendahl Informationstechnik**
Freelancer (student job): Software development

10/2008 - 09/2011 — **Westmuensterhaendler GbR - www.Holzschuhe.de**
Partner (student job): Distribution of traditional German wooden shoes

05/2005 - 05/2008	**GWAVA EMEA GmbH** Software developer (student job) 06/2006 - 09/2006: Full-time developer at GWAVA Inc., Montreal, Canada

Education

Since 04/2012	**Fraunhofer Institute for Computer Graphics Research IGD** Ph.D.: Capacitive Sensing and Communication for Ubiquitous Interaction and Perception
03/2010 - 03/2012	**Information System Technology, Technische Universität Darmstadt** Master of Science: Embedded systems, Human-Computer Interaction, Ubiquitous Computing (thesis: 1.0, overall: 1.3)
10/2006 - 03/2010	**Information System Technology, Technische Universität Darmstadt** Bachelor of Science: Embedded systems, Human-Computer Interaction, Ubiquitous Computing (thesis: 1.0, overall: 2.1)
08/2005 - 05/2006	**Retirement Home 'Holthues Hoff'** Community service: Physical and social care (initiated weekly cooking groups, bowling, organization of events)
06/1996 - 06/2005	**Alexander-Hegius-Gymnasium Ahaus** Abitur 08/2003 - 12/2003: Budehaven Community School, Bude, England

Professional Services

Organizer: Workshop on Smart Objects 2015

Program Committee: iWOAR 2015

Reviewer: Augmented Human 2015, CHI 2015 / 2013, IUI 2015, UbiComp 2015

Thursday 28$^{\text{th}}$ May, 2015

Bibliography

[AB10] ALTUN K., BARSHAN B.: Human activity recognition using inertial/magnetic sensor units. In *Human Behavior Understanding*, Salah A., Gevers T., Sebe N., Vinciarelli A., (Eds.), vol. 6219 of *Lecture Notes in Computer Science*. Springer Berlin / Heidelberg, 2010, pp. 38–51. 104

[AD13] ALLHOFF K. T., DROSSEL B.: When do evolutionary food web models generate complex networks? *Journal of Theoretical Biology 334*, 0 (2013), 122 – 129. 140

[AJT05] AMFT O., JUNKER H., TRÖSTER G.: Detection of eating and drinking arm gestures using inertial body-worn sensors. In *Proceedings of the 9th IEEE International Symposium on Wearable Computers (ISWC '05)* (2005), IEEE, pp. 160–163. 29, 37, 38

[AKKM14] ADIB F., KABELAC Z., KATABI D., MILLER R. C.: 3d tracking via body radio reflections. In *Proceedings of the 11th USENIX Conference on Networked Systems Design and Implementation* (Berkeley, CA, USA, 2014), NSDI'14, USENIX Association, pp. 317–329. 32, 37, 38, 135

[AM03] AITENBICHLER E., MUHLHAUSER M.: An ir local positioning system for smart items and devices. In *Distributed Computing Systems Workshops, 2003. Proceedings. 23rd International Conference on* (May 2003), pp. 334–339. 38

[AMGC02] ARULAMPALAM M. S., MASKELL S., GORDON N., CLAPP T.: A tutorial on particle filters for online nonlinear/non-Gaussian Bayesian tracking. *IEEE Transactions on Signal Processing* (2002). 89

[ASLT05] AMFT O., STÄGER M., LUKOWICZ P., TRÖSTER G.: Analysis of chewing sounds for dietary monitoring. In *Proceedings of the 7th International Conference on Ubiquitous Computing* (Berlin, Heidelberg, 2005), UbiComp'05, Springer-Verlag, pp. 56–72. 34, 76, 104

[ATB11] AURELIEN TABARD JUAN-DAVID HINCAPIE-RAMOS M. E., BARDRAM J. E.: The eLabBench: an interactive tabletop system for the biology laboratory. In *ITS '11* (2011), pp. 202–211. 114

[AW09] AARTS E., WICHERT R.: Ambient intelligence. In *Technology Guide*. Springer, 2009, pp. 244–249. 7

[Bax96] BAXTER L.: *Capacitive Sensors: Design and Applications*. IEEE Press Series on Electronics Technology. John Wiley & Sons, 1996. xiii, 11, 12, 21, 80

[BBRS06] BALLAGAS R., BORCHERS J., ROHS M., SHERIDAN J. G.: The smart phone: a ubiquitous input device. *IEEE Pervasive Computing 5*, 1 (2006), 70–77. 125

[BCS*12] BASIRICÒ L., COSSEDDU P., SCIDÀ A., FRABONI B., MALLIARAS G., BONFIGLIO A.: Electrical characteristics of ink-jet printed, all-polymer electrochemical transistors. *Organic Electronics 13*, 2 (2012), 244–248. 50

[Bec13] BECK S.: *A Gesture Recognition Device with Visual Feedback based on Capacitive Proximity Sensing*. Bachelor's thesis, Technische Universität Darmstadt, 2013. 124

[Ber12] BERGHÖFER Y.: *Human-Machine-Interfaces in Automotive Environments using Capac-*

itive Proximity Sensors. Master's thesis, Technische Universität Darmstadt, 2012. xiii, xv, xvii, xviii, 5, 19, 44, 50, 57, 83, 86, 91, 80, 97, 98, 99, 100, 101

[BGC09] BREZMES T., GORRICHO J.-L., COTRINA J.: Activity recognition from accelerometer data on a mobile phone. In *Distributed Computing, Artificial Intelligence, Bioinformatics, Soft Computing, and Ambient Assisted Living*, Omatu S., Rocha M., Bravo J., Fernández F., Corchado E., Bustillo A., Corchado J., (Eds.), vol. 5518 of *Lecture Notes in Computer Science*. Springer Berlin / Heidelberg, 2009, pp. 796–799. 35, 38, 104

[BH09] BRAUN A., HAMISU P.: Using the human body field as a medium for natural interaction. In *Proceedings of the 2Nd International Conference on PErvasive Technologies Related to Assistive Environments* (New York, NY, USA, 2009), PETRA '09, ACM, pp. 50:1–50:7. 16, 55

[BH11] BRAUN A., HAMISU P.: Designing a multi-purpose capacitive proximity sensing input device. *PETRA '11* (2011), 151–158. 80

[BHH*13] BRÄNZEL A., HOLZ C., HOFFMANN D., SCHMIDT D., KNAUST M., LÜHNE P., MEUSEL R., RICHTER S., BAUDISCH P.: Gravityspace: Tracking users and their poses in a smart room using a pressure-sensing floor. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems* (New York, NY, USA, 2013), CHI '13, ACM, pp. 725–734. http://doi.acm.org/10.1145/2470654.2470757. xiv, 27, 28

[BHW12] BRAUN A., HEGGEN H., WICHERT R.: Capfloor – a flexible capacitive indoor localization system. In *Evaluating AAL Systems Through Competitive Benchmarking. Indoor Localization and Tracking*, Chessa S., Knauth S., (Eds.), vol. 309 of *Communications in Computer and Information Science*. Springer Berlin Heidelberg, 2012, pp. 26–35. 23, 24, 38

[BI04] BAO L., INTILLE S.: Activity recognition from user-annotated acceleration data. In *Pervasive Computing*, Ferscha A., Mattern F., (Eds.), vol. 3001 of *Lecture Notes in Computer Science*. Springer Berlin Heidelberg, 2004, pp. 1–17. 34, 37, 38, 103, 104

[BIH08] BUTLER A., IZADI S., HODGES S.: Sidesight: Multi-"touch" interaction around small devices. In *Proceedings of the 21st Annual ACM Symposium on User Interface Software and Technology* (New York, NY, USA, 2008), UIST '08, ACM, pp. 201–204. 33

[Bis06] BISHOP C. M.: *Pattern Recognition and Machine Learning (Information Science and Statistics)*. Springer-Verlag New York, Inc., 2006. 89

[BJK*07] BEETZ M., JAN B., KIRSCH R., MALDONADO A., MÜLLER A., RUSU R. B.: The assistive kitchen — a demonstration scenario for cognitive technical systems. In *in Proceedings of the 4th COE Workshop on Human Adaptive Mechatronics* (2007). 114

[BLL06] BABINEAU D., LONGTIN A., LEWIS J. E.: Modeling the electric field of weakly electric fish. *Journal of Experimental Biology 209*, 18 (2006), 3636–3651. 8

[BLO*05] BORRIELLO G., LIU A., OFFER T., PALISTRANT C., SHARP R.: Walrus: Wireless acoustic location with room-level resolution using ultrasound. In *Proceedings of the 3rd International Conference on Mobile Systems, Applications, and Services* (New York, NY, USA, 2005), MobiSys '05, ACM, pp. 191–203. 30

[BLvLS10] BERLIN E., LIU J., VAN LAERHOVEN K., SCHIELE B.: Coming to grips with the objects we grasp: Detecting interactions with efficient wrist-worn sensors. In *Proceedings of the Fourth International Conference on Tangible, Embedded, and Embodied Interaction* (New York, NY, USA, 2010), TEI '10, ACM, pp. 57–64. 34, 38, 39, 104

[BM86] BUXTON B., MYERS B. A.: A study in two-handed input. In *CHI '86* (1986), pp. 321–

326. 93

[BNS*14] BRAUN A., NEUMANN S., SCHMIDT S., WICHERT R., KUIJPER A.: Towards interactive car interiors: the active armrest. In *Proceedings of the 8th Nordic Conference on Human-Computer Interaction: Fun, Fast, Foundational* (2014), ACM, pp. 911–914. 57

[BO10] BARRETT G., OMOTE R.: Projected-capacitive touch technology. *Information Display 26*, 3 (2010), 16–21. 15, 21

[Bra09] BRAUN A.: *Modeling of Input Devices for Natural Interaction from Low-Level Capacitive Proximity Sensor Data*. Master's thesis, Technische Universität Darmstadt, 2009. 16, 24, 38

[Bra14] BRAUN A.: *Application and validation of capacitive proximity sensing systems in smart environments*. PhD thesis, Technische Universität Darmstadt, 2014. 7, 58

[Bre02] BRENT R.: Algorithms for Minimization Without Derivatives. In *Dover Publications* (2002). 86

[BSB11] B. SEEWALD M. JOHN J. S., BELBACHIR A. N.: Silvergame – a project aimed at social integration and multimedia interaction for the elderly, 2011. `http://www.silvergame.eu/assets/Silvergame_01_EN.pdf` (accessed 30/09/2014). 27

[Bux96] BUXTON W.: Living in augmented reality: Ubiquitous media and reactive environments. In *Video Mediated Communication*, Finn K., Sellen A., Wilber S., (Eds.). Erlbaum, Hillsdale, NJ, 1996. 43

[BV12] BORAZIO M., VAN LAERHOVEN K.: Combining wearable and environmental sensing into an unobtrusive tool for long-term sleep studies. In *Proceedings of the 2nd ACM SIGHIT symposium on International health informatics - IHI '12* (2012), ACM Press, pp. 71–80. 76, 104

[BZWK*14] BRAUN A., ZANDER-WALZ S., KREPP S., RUS S., WICHERT R., KUIJPER A.: Captap - combining capacitive gesture recognition and knock detection. In *Working Paper* (2014). xv, 55, 57

[CABL13] CHENG J., AMFT O., BAHLE G., LUKOWICZ P.: Designing sensitive wearable capacitive sensors for activity recognition. *Sensors Journal, IEEE 13*, 10 (Oct 2013), 3935–3947. xiv, 22

[CAL10] CHENG J., AMFT O., LUKOWICZ P.: Active capacitive sensing: Exploring a new wearable sensing modality for activity recognition. In *Proceedings of the 8th International Conference on Pervasive Computing* (Berlin, Heidelberg, 2010), Pervasive'10, Springer-Verlag, pp. 319–336. vi, 3, 22, 38, 39, 70, 104

[CATC14] CONSTANDACHE I., AGARWAL S., TASHEV I., CHOUDHURY R. R.: Daredevil: Indoor location using sound. *SIGMOBILE Mob. Comput. Commun. Rev. 18*, 2 (June 2014), 9–19. 29, 38

[CBL08] CHENG J., BANNACH D., LUKOWICZ P.: On body capacitive sensing for a simple touchless user interface. In *Medical Devices and Biosensors, 2008. ISSS-MDBS 2008. 5th International Summer School and Symposium on* (June 2008), pp. 113–116. xiv, 21, 22, 38

[CCGP13] CHEN K.-Y., COHN G. A., GUPTA S., PATEL S. N.: utouch: Sensing touch gestures on unmodified lcds. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems* (New York, NY, USA, 2013), CHI '13, ACM, pp. 2581–2584. 38

[CGL*12] COHN G., GUPTA S., LEE T.-J., MORRIS D., SMITH J. R., REYNOLDS M. S., TAN D. S., PATEL S. N.: An ultra-low-power human body motion sensor using static electric field sensing. In *Proceedings of the 2012 ACM Conference on Ubiquitous Computing* (New York, NY, USA, 2012), UbiComp '12, ACM, pp. 99–102. vi, 38, 39, 70, 125

[CJC*06] CRISPIN X., JAKOBSSON F. L. E., CRISPIN A., GRIM P. C. M., ANDERSSON P., VOLODIN A., HAESENDONCK C. V., AUWERAER M. V. D., SALANECK W. R., BERGGREN M.: The origin of the high conductivity of (pedot-pss) plastic electrodes. *Chemistry of Materials 18*, 18 (2006), 4354–4360. 50, 70

[CLC*06] CHANG K.-H., LIU S.-Y., CHU H.-H., HSU J.-J., CHEN C., LIN T.-Y., CHEN C.-Y., HUANG P.: The diet-aware dining table: Observing dietary behaviors over a tabletop surface. In *Pervasive '06*. 2006, pp. 366–382. 114

[CLC*10] CHENG K.-Y., LIANG R.-H., CHEN B.-Y., LAING R.-H., KUO S.-Y.: icon: Utilizing everyday objects as additional, auxiliary and instant tabletop controllers. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems* (New York, NY, USA, 2010), CHI '10, ACM, pp. 1155–1164. 28

[CMPT12] COHN G., MORRIS D., PATEL S., TAN D.: Humantenna: using the body as an antenna for real-time whole-body interaction. In *CHI '12* (2012), pp. 1901–1910. 3, 24, 38, 39, 70, 79, 137

[Cor01] CORRADINI A.: Dynamic time warping for off-line recognition of a small gesture vocabulary. In *Proceedings of IEEE ICCV Workshop on Recognition, Analysis, and Tracking of Faces and Gestures in Real-Time Systems* (2001), pp. 82–89. 90

[CPIP10] CHINTALAPUDI K., PADMANABHA IYER A., PADMANABHAN V. N.: Indoor localization without the pain. In *Proceedings of the Sixteenth Annual International Conference on Mobile Computing and Networking* (New York, NY, USA, 2010), MobiCom '10, ACM, pp. 173–184. 32, 38

[CPL12] CHEN L., PAN G., LI S.: Touch-driven interaction via an NFC-enabled smartphone. In *PerCom '12 Workshops* (2012), pp. 504–506. 38, 71, 136

[Cre07] CREMER M.: Ueber die Registrierung mechanischer Vorgänge auf electrischem Wege, speziell mit Hilfe des Saitengalvanometers und Saitenelectrometers. *Münchner Medizinische Wochenschrift 54* (August 1907), 1629–1630. 9

[CW05] CLAUSERT H., WIESEMANN G.: *Grundgebiete der Elektrotechnik 1*. No. Bd. 1. Oldenbourg, 2005. 10, 11, 12

[Cyp12] CYPRESS SEMICONDUCTOR CORP.: Cypress TrueTouch Touchscreen Solution Drives "Floating Touch" Navigation Feature in New Xperia™ sola Smartphone from Sony Mobile Communications, 2012. *Press release, http://www.cypress.com/?rID=60561 (accessed 12/17/2012)*. 79, 97

[DBM14] DJAKOW M., BRAUN A., MARINC A.: Movibed-sleep analysis using capacitive sensors. In *Universal Access in Human-Computer Interaction. Design for All and Accessibility Practice*. Springer, 2014, pp. 171–181. 57

[DKA14] DUONG M. Q., KAWAHARA Y., ASAMI T.: Design of touch-sensitive surface with arbitrary shape based on time-domain reflectometry using injket printing. In *Radio and Wireless Symposium (RWS), 2014 IEEE* (Jan 2014), pp. 16–18. 31, 38

[DL01] DIETZ P., LEIGH D.: Diamondtouch: A multi-user touch technology. In *UIST '01* (2001), pp. 219–226. 17, 70, 136

[DL12] DODSON B., LAM M.: Micro-interactions with nfc-enabled mobile phones. In *Mobile Computing, Applications, and Services*, vol. 95. Springer, 2012, pp. 118–136. 71

[ECC13] EVANS D., CLAIBORNE J., CURRIE S.: *The Physiology of Fishes, Fourth Edition*. CRC

Marine Biology Series. Taylor & Francis, 2013. 8, 9

[FBL14] FREEMAN E., BREWSTER S., LANTZ V.: Illuminating Gesture Interfaces with Interactive Light Feedback. In *Beyond the Switch: Explicit and Implicit Interaction with Light Workshop* (2014). xx, 131, 132

[FGPK14] FU B., GROSSE-PUPPENDAHL T., KUIJPER A.: A gesture recognition method for proximity-sensing surfaces in smart environments. In *Working Paper* (2014). 92

[Fia05] FIALA M.: Artag, a fiducial marker system using digital techniques. In *Computer Vision and Pattern Recognition, 2005. CVPR 2005. IEEE Computer Society Conference on* (June 2005), vol. 2, pp. 590–596 vol. 2. 28, 38

[FL96] FOSTER K. R., LUKASKI H. C.: Whole-body impedance–what does it measure? *The American journal of clinical nutrition 64*, 3 (1996), 388S–396S. xiii, 18, 37, 70

[FPR05] FISHKIN K., PHILIPOSE M., REA A.: Hands-on rfid: wireless wearables for detecting use of objects. In *Proceedings of the 9th IEEE International Symposium on Wearable Computers (ISWC '05)* (Oct. 2005), pp. 38–41. 104

[FWK*13] FEI S., WEBB A. M., KERNE A., QU Y., JAIN A.: Peripheral Array of Tangible NFC Tags: Positioning Portals for Embodied Trans-Surface Interaction. In *ITS '13* (2013). 38, 71

[Gat06] GATES E. D.: *Introduction to Electronics*. Thomson Learning, 2006. 47

[GF01] GREENBERG S., FITCHETT C.: Phidgets: easy development of physical interfaces through physical widgets. In *UIST '01* (2001), pp. 209–218. 44

[Gli00] GLINSKY A.: *Theremin: Ether Music and Espionage*. Music in American life. University of Illinois Press, 2000. 9, 10, 15, 22

[GLIA*14] GOEL M., LEE B., ISLAM AUMI M. T., PATEL S., BORRIELLO G., HIBINO S., BEGOLE B.: Surfacelink: Using inertial and acoustic sensing to enable multi-device interaction on a surface. In *Proceedings of the 32Nd Annual ACM Conference on Human Factors in Computing Systems* (New York, NY, USA, 2014), CHI '14, ACM, pp. 1387–1396. 29

[GMB*11] GREENBERG S., MARQUARDT N., BALLENDAT T., DIAZ-MARINO R., WANG M.: Proxemic interactions: The new ubicomp? *interactions 18*, 1 (Jan. 2011), 42–50. http://doi.acm.org/10.1145/1897239.1897250. xiii, xiv, xv, xvii, xx, 2, 8, 43, 44, 59, 78, 79, 133

[GMPT12] GUPTA S., MORRIS D., PATEL S., TAN D.: Soundwave: Using the doppler effect to sense gestures. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems* (New York, NY, USA, 2012), CHI '12, ACM, pp. 1911–1914. 30, 37, 38, 97, 135

[Goe58] GOERTZEL G.: An algorithm for the evaluation of finite trigonometric series. *American Mathematical Monthly 65*, 1 (1958), 34–35. 65

[GP12] GROSSE-PUPPENDAHL T.: *Multi-hand interaction using custom capacitive proximity sensors*. Master's thesis, Technische Universität Darmstadt, 2012. xvii, 20, 79, 89, 80, 93, 94

[GPB12] GROSSE-PUPPENDAHL T., BRAUN A.: Honeyfish - a high resolution gesture recognition system based on capacitive proximity sensing. In *Embedded World Conference 2012* (2012), pp. 1–10. vii, 16, 80, 92

[GPBB12] GROSSE-PUPPENDAHL T., BERLIN E., BORAZIO M.: Enhancing accelerometer-based activity recognition with capacitive proximity sensing. In *AmI '12* (2012). (publication date 11/2012). vii, 5, 38, 39, 57, 104, 135

[GPBB*13] GROSSE-PUPPENDAHL T., BERGHOEFER Y., BRAUN A., WIMMER R., KUIJPER A.: OpenCapSense: A Rapid Prototyping Toolkit for Pervasive Interaction using Capacitive Sensing. In *PerCom '13* (2013), pp. 152–159. vi, 4, 38, 44, 57, 61, 66, 70, 115, 119, 121, 135

[GPBK*13] GROSSE-PUPPENDAHL T., BENCHEA S., KAMIETH F., BRAUN A., SCHUSTER C.: Unobtrusive recognition of working situations. In *Distributed, Ambient, and Pervasive Interactions*, Streitz N., Stephanidis C., (Eds.), vol. 8028 of *Lecture Notes in Computer Science*. Springer Berlin Heidelberg, 2013, pp. 115–121. vii, 5, 57, 104, 114, 135

[GPBKK13] GROSSE-PUPPENDAHL T., BRAUN A., KAMIETH F., KUIJPER A.: Swiss-Cheese Extended: An Object Recognition Method for Ubiquitous Interfaces based on Capacitive Proximity Sensing. In *CHI '13* (2013), pp. 1401–1410. vii, 5, 37, 70, 80, 135

[GPBW14a] GROSSE-PUPPENDAHL T., BECK S., WILBERS D.: Rainbowfish: Visual feedback on gesture-recognizing surfaces. In *CHI '14 Extended Abstracts on Human Factors in Computing Systems* (New York, NY, USA, 2014), CHI EA '14, ACM, pp. 427–430. viii, 5, 124, 135, 136

[GPBW*14b] GROSSE-PUPPENDAHL T., BECK S., WILBERS D., ZEISS S., VON WILMSDORFF J., KUIJPER A.: Ambient gesture-recognizing surfaces with visual feedback. In *Distributed, Ambient, and Pervasive Interactions*, Streitz N., Markopoulos P., (Eds.), vol. 8530 of *Lecture Notes in Computer Science*. Springer International Publishing, 2014, pp. 97–108. viii, 5, 124, 135, 136

[GPHW*14] GROSSE-PUPPENDAHL T., HERBER S., WIMMER R., ENGLERT F., BECK S., VON WILMSDORFF J., WICHERT R., KUIJPER A.: Capacitive near-field communication for ubiquitous interaction and perception. In *Proceedings of the 2014 ACM International Joint Conference on Pervasive and Ubiquitous Computing* (New York, NY, USA, 2014), UbiComp '14 Adjunct, ACM, pp. 231–242. 4, 5, 37, 38, 40, 59, 135, 136, 137

[GPMB11] GROSSE-PUPPENDAHL T., MARINC A., BRAUN A.: Classification of User Postures with Capacitive Proximity Sensors in AAL-Environments. In *Ambient Intelligence* (2011), Keyson D. V., Maher M. L., Streitz N., Cheok A., Augusto J. C., Wichert R., Englebienne G., Aghajan H., Kröse B. J. A., (Eds.), vol. 7040 of *Lecture Notes in Computer Science*, Springer, pp. 314–323. vii, 5, 38, 47, 57, 104, 114, 135

[GRP10] GUPTA S., REYNOLDS M. S., PATEL S. N.: Electrisense: Single-point sensing using emi for electrical event detection and classification in the home. In *Proceedings of the 12th ACM International Conference on Ubiquitous Computing* (New York, NY, USA, 2010), Ubicomp '10, ACM, pp. 139–148. 25

[GSO*14] GONG N.-W., STEIMLE J., OLBERDING S., HODGES S., GILLIAN N. E., KAWAHARA Y., PARADISO J. A.: Printsense: A versatile sensing technique to support multimodal flexible surface interaction. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems* (New York, NY, USA, 2014), CHI '14, ACM, pp. 1407–1410. 21, 35, 37, 38

[Hal66] HALL E.: *The Hidden Dimension*. A Doubleday anchor book. Anchor Books, 1966. 2

[Han05] HAN J. Y.: Low-cost multi-touch sensing through frustrated total internal reflection. In *Proceedings of the 18th Annual ACM Symposium on User Interface Software and Technology* (New York, NY, USA, 2005), UIST '05, ACM, pp. 115–118. 26, 28, 33, 37, 38

[HB13] HOLZ C., BAUDISCH P.: Fiberio: A touchscreen that senses fingerprints. In *Proceedings of the 26th Annual ACM Symposium on User Interface Software and Technology* (New York, NY, USA, 2013), UIST '13, ACM, pp. 41–50. http://doi.acm.org/10.1145/2501988.2502021. xiv, 26, 27, 37, 38

[HBS07] HUYNH T., BLANKE U., SCHIELE B.: Scalable recognition of daily activities with wearable sensors. In *Location- and Context-Awareness*, Hightower J., Schiele B., Strang T., (Eds.), vol. 4718 of *Lecture Notes in Computer Science*. Springer Berlin Heidelberg, 2007, pp. 50–67. 35

[HBW12] HWANG S., BIANCHI A., WOHN K.: Micpen: Pressure-sensitive pen interaction using microphone with standard touchscreen. In *CHI '12 Extended Abstracts on Human Factors in Computing Systems* (New York, NY, USA, 2012), CHI EA '12, ACM, pp. 1847–1852. 37

[hed] HedgeHog activity logger. `http://www.ess.tu-darmstadt.de/hedgehog`. Accessed: 17-06-2012. 105

[Hei77] HEILIGENBERG W.: *Principles of electrolocation and jamming avoidance in electric fish: a neuroethological approach*. Studies of brain function. Springer-Verlag, 1977. 9

[Her13] HERBER S.: *Tangible Interaction using Capacitive Near-Field Communication*. Bachelor's thesis, Technische Universität Darmstadt, 2013. 60

[HH01] HOWARD B., HOWARD S.: Lightglove: Wrist-worn virtual typing and pointing. In *Proceedings of the 5th IEEE International Symposium on Wearable Computers* (Washington, DC, USA, 2001), ISWC '01, IEEE Computer Society, pp. 172–. 33, 38

[HH08] HARRISON C., HUDSON S. E.: Scratch input: Creating large, inexpensive, unpowered and mobile finger input surfaces. In *Proceedings of the 21st Annual ACM Symposium on User Interface Software and Technology* (New York, NY, USA, 2008), UIST '08, ACM, pp. 205–208. 28, 29, 34, 37, 38

[HHL09] HUANG C.-F., HUNG Y.-C., LIU C.-L.: P-53: Precise location of touch panel by employing the time-domain reflectometry. *SID Symposium Digest of Technical Papers 40*, 1 (2009), 1291–1293. 31

[HIB*07] HODGES S., IZADI S., BUTLER A., RRUSTEMI A., BUXTON B.: Thinsight: Versatile multi-touch sensing for thin form-factor displays. In *Proceedings of the 20th Annual ACM Symposium on User Interface Software and Technology* (New York, NY, USA, 2007), UIST '07, ACM, pp. 259–268. 136

[Hop81] HOPKINS C. D.: The neuroethology of electric communication. *Trends in Neurosciences 4*, 0 (1981), 4 – 6. 8, 9

[HPC14] HUGHES D., PROFITA H., CORRELL N.: Switchback: An on-body rf-based gesture input device. In *Proceedings of the 2014 ACM International Symposium on Wearable Computers* (New York, NY, USA, 2014), ISWC '14, ACM, pp. 63–66. 31, 37, 38

[HPSH00] HINCKLEY K., PIERCE J., SINCLAIR M., HORVITZ E.: Sensing techniques for mobile interaction. In *Proceedings of the 13th Annual ACM Symposium on User Interface Software and Technology* (New York, NY, USA, 2000), UIST '00, ACM, pp. 91–100. 33

[HSAT10] HOLLECZEK T., SCHOCH J., ARNRICH B., TRÖANDSTER G.: Recognizing turns and other snowboarding activities with a gyroscope. In *Proceedings of the 14th IEEE International Symposium on Wearable Computers (ISWC '10)* (Oct. 2010), pp. 1–8. 104

[HSP12] HARRISON C., SATO M., POUPYREV I.: Capacitive fingerprinting: Exploring user differentiation by sensing electrical properties of the human body. In *Proceedings of the 25th Annual ACM Symposium on User Interface Software and Technology* (New York, NY, USA, 2012), UIST '12, ACM, pp. 537–544. 18, 59, 70

[HTM10] HARRISON C., TAN D., MORRIS D.: Skinput: Appropriating the body as an input surface. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems* (New York, NY, USA, 2010), CHI '10, ACM, pp. 453–462. 28, 29, 36, 37, 38

[HV11] HOLMAN D., VERTEGAAL R.: Tactiletape: Low-cost touch sensing on curved surfaces. In *Proceedings of the 24th Annual ACM Symposium Adjunct on User Interface Software and Technology* (New York, NY, USA, 2011), UIST '11 Adjunct, ACM, pp. 17–18. 21

[HVC*14] HODGES S., VILLAR N., CHEN N., CHUGH T., QI J., NOWACKA D., KAWAHARA Y.: Circuit stickers: Peel-and-stick construction of interactive electronic prototypes. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems* (New York, NY, USA, 2014), CHI '14, ACM, pp. 1743–1746. 21

[HWL*03] HELAL S., WINKLER B., LEE C., KADDOURA Y., RAN L., GIRALDO C., KUCHIBHOTLA S., MANN W.: Enabling location-aware pervasive computing applications for the elderly. In *Pervasive Computing and Communications, 2003. (PerCom 2003). Proceedings of the First IEEE International Conference on* (March 2003), pp. 531–536. 30, 38

[IB98] ISARD M., BLAKE A.: Condensation | conditional density propagation for visual tracking. *International Journal of Computer Vision* (1998), 5–28. 88

[IHB*07] IZADI S., HODGES S., BUTLER A., RRUSTEMI A., BUXTON B.: Thinsight: Integrated optical multi-touch sensing through thin form-factor displays. In *Proceedings of the 2007 Workshop on Emerging Displays Technologies: Images and Beyond: The Future of Displays and Interacton* (New York, NY, USA, 2007), EDT '07, ACM. 33

[Inv14] INVENSENSE INC.: MPU-9255 Product Specification , 2014. http://invensense.com/mems/gyro/documents/PS-MPU-9255.pdf (date accessed: 2014/10/15). 34

[IWO*99] ISHII H., WISNESKI C., ORBANES J., CHUN B., PARADISO J.: Pingpongplus: Design of an athletic-tangible interface for computer-supported cooperative play. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems* (New York, NY, USA, 1999), CHI '99, ACM, pp. 394–401. 29, 38

[JGAK07] JORDÀ S., GEIGER G., ALONSO M., KALTENBRUNNER M.: The reactable: Exploring the synergy between live music performance and tabletop tangible interfaces. In *Proceedings of the 1st International Conference on Tangible and Embedded Interaction* (New York, NY, USA, 2007), TEI '07, ACM, pp. 139–146. 136

[Jor10] JORDÀ S.: The reactable: Tangible and tabletop music performance. In *CHI '10 Extended Abstracts on Human Factors in Computing Systems* (New York, NY, USA, 2010), CHI EA '10, ACM, pp. 2989–2994. 28, 38

[JS09] JINHA SEONG WOOHUN LEE Y.-K. L.: Why we cannot work without paper even in a computerized work environment. In *CHI '09 EA* (2009), pp. 4105–4110. 114

[KB07] KALTENBRUNNER M., BENCINA R.: reactivision: A computer-vision framework for table-based tangible interaction. In *Proceedings of the 1st International Conference on Tangible and Embedded Interaction* (New York, NY, USA, 2007), TEI '07, ACM, pp. 69–74. 28, 38

[KBHG05] KROHN A., BEIGL M., HAZAS M., GELLERSEN H.-W.: Using fine-grained infrared positioning to support the surface-based activities of mobile users. In *Distributed Computing Systems Workshops, 2005. 25th IEEE International Conference on* (June 2005), pp. 463–468. 33, 38

[KFB*97] KURTENBACH G., FITZMAURICE G., BAUDEL T., BUXTON B., EAST R. S.: The Design of a GUI Paradigm based on Tablets, Two-hands, and Transparency. In *CHI '97* (1997), pp. 35–42. 92, 93

[KHC*13] KAWAHARA Y., HODGES S., COOK B. S., ZHANG C., ABOWD G. D.: Instant inkjet circuits: Lab-based inkjet printing to support rapid prototyping of ubicomp devices. In *Proceedings of the 2013 ACM International Joint Conference on Pervasive and Ubiquitous Computing* (New York, NY, USA, 2013), UbiComp '13, ACM, pp. 363–372. 21, 66, 70

[KHI*12] KIM D., HILLIGES O., IZADI S., BUTLER A. D., CHEN J., OIKONOMIDIS I., OLIVIER P.: Digits: Freehand 3d interactions anywhere using a wrist-worn gloveless sensor. In *Proceedings of the 25th Annual ACM Symposium on User Interface Software and Technology* (New York, NY, USA, 2012), UIST '12, ACM, pp. 167–176. 38

[Kir14] KIRCHBUCHNER F.: *User Tracking and Behavior Recognition Based on a Capacitive Indoor Localization System.* Master's thesis, Technische Universität Darmstadt, 2014. 3

[KL08] KUNZE K., LUKOWICZ P.: Dealing with sensor displacement in motion-based onbody activity recognition systems. In *Proceedings of the 10th International Conference on Ubiquitous Computing* (New York, NY, USA, 2008), UbiComp '08, ACM, pp. 20–29. 37

[KMA01] KOLLER-MEIER E. B., ADE F.: Tracking multiple objects using the Condensation algorithm. *Robotics and Autonomous Systems 34*, 2-3 (2001), 93–105. 88

[KR09] KALGAONKAR K., RAJ B.: One-handed gesture recognition using ultrasonic doppler sonar. In *Acoustics, Speech and Signal Processing, 2009. ICASSP 2009. IEEE International Conference on* (April 2009), pp. 1889–1892. 30, 38

[KSK*10] KIVIKUNNAS S., STRÖMMER E., KORKALAINEN M., HEIKKILÄ T., HAVERINEN M.: Intelligent furniture and their ubiquitous use scenarios. In *Proceedings of the AALIANCE conference* (2010). 114

[KSSF03] KRAUSE A., SIEWIOREK D., SMAILAGIC A., FARRINGDON J.: Unsupervised, dynamic identification of physiological and activity context in wearable computing. In *Proceedings of the 7th IEEE International Symposium on Wearable Computers (ISWC '03)* (Washington, DC, USA, Oct. 2003), IEEE Computer Society, pp. 88–97. 104

[KTG14] KELLOGG B., TALLA V., GOLLAKOTA S.: Bringing gesture recognition to all devices. In *Proceedings of the 11th USENIX Conference on Networked Systems Design and Implementation* (Berkeley, CA, USA, 2014), NSDI'14, USENIX Association, pp. 303–316. 32

[Lan05] LANDT J.: The history of rfid. *Potentials, IEEE 24*, 4 (Oct 2005), 8–11. 34

[LB13] LOPES P., BAUDISCH P.: Muscle-propelled force feedback: Bringing force feedback to mobile devices. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems* (New York, NY, USA, 2013), CHI '13, ACM, pp. 2577–2580. 125, 136

[LCC*13] LIANG R.-H., CHENG K.-Y., CHAN L., PENG C.-X., CHEN M. Y., LIANG R.-H., YANG D.-N., CHEN B.-Y.: Gaussbits: Magnetic tangible bits for portable and occlusion-free near-surface interactions. In *Proceedings of the SIGCHI Conference on*

Human Factors in Computing Systems (New York, NY, USA, 2013), CHI '13, ACM, pp. 1391–1400. 136

[Lea14] LEAP MOTION, INC: Leap Motion, 2014. http://www.leapmotion.com (date accessed: 2014/10/12). 26, 38, 97

[LGTIK14] LE GOC M., TAYLOR S., IZADI S., KESKIN C.: A low-cost transparent electric field sensor for 3d interaction on mobile devices. In *Proceedings of the 32Nd Annual ACM Conference on Human Factors in Computing Systems* (New York, NY, USA, 2014), CHI '14, ACM, pp. 3167–3170. http://doi.acm.org/10.1145/2556288.2557331. xiv, 23, 24

[LJA*12] LIU J., JOHNS E., ATALLAH L., PETTITT C., LO B., FROST G., YANG G.-Z.: An intelligent food-intake monitoring system using wearable sensors. In *BSN '12* (2012), IEEE, pp. 154–160. 76

[LJJ11] LOPES P., JOTA R., JORGE J. A.: Augmenting touch interaction through acoustic sensing. In *Proceedings of the ACM International Conference on Interactive Tabletops and Surfaces* (New York, NY, USA, 2011), ITS '11, ACM, pp. 53–56. 29

[LLS11] LEE S., LI B., STARNER T.: Airtouch: Synchronizing in-air hand gesture and on-body tactile feedback to augment mobile gesture interaction. In *Wearable Computers (ISWC), 2011 15th Annual International Symposium on* (June 2011), pp. 3–10. 33, 38

[LP09] LU Y., PAYANDEH S.: Intelligent cooperative tracking in multi-camera systems. In *Intelligent Systems Design and Applications, 2009. ISDA '09. Ninth International Conference on* (Nov 2009), pp. 608–613. 27

[LRF12] LEI J., REN X., FOX D.: Fine-grained kitchen activity recognition using rgb-d. In *Proceedings of the 2012 ACM Conference on Ubiquitous Computing* (New York, NY, USA, 2012), UbiComp '12, ACM, pp. 208–211. 28

[MBBN05] MIDDLETON L., BUSS A., BAZIN A., NIXON M.: A floor sensor system for gait recognition. In *Automatic Identification Advanced Technologies, 2005. Fourth IEEE Workshop on* (Oct 2005), pp. 171–176. 35

[MBMK13] MAJEWSKI M., BRAUN A., MARINC A., KUIJPER A.: Providing visual support for selecting reactive elements in intelligent environments. *Transactions on Computational Science XVIII 7848* (2013), 248–263. 125

[MBRS14] MENNICKEN S., BRUSH A. J. B., ROSEWAY A., SCOTT J.: Exploring interactive furniture with emotocouch. In *Adjunct Proceedings of ACM UbiComp 2014 (videos track)* (September 2014), ACM. 136

[MDL*02] MCELLIGOTT L., DILLON M., LEYDON K., RICHARDSON B., FERNSTRÖM M., PARADISO J.: 'forse fields' - force sensors for interactive environments. In *UbiComp 2002: Ubiquitous Computing*, Borriello G., Holmquist L., (Eds.), vol. 2498 of *Lecture Notes in Computer Science*. Springer Berlin Heidelberg, 2002, pp. 168–175. 35

[MG09] MOLYNEAUX D., GELLERSEN H.: Projected interfaces: Enabling serendipitous interaction with smart tangible objects. In *Proceedings of the 3rd International Conference on Tangible and Embedded Interaction* (New York, NY, USA, 2009), TEI '09, ACM, pp. 385–392. 28, 38

[MGS08] MOLYNEAUX D., GELLERSEN H., SCHIELE B.: Vision-based detection of mobile smart objects. In *Smart Sensing and Context*, Roggen D., Lombriser C., Tröster G., Kortuem G., Havinga P., (Eds.), vol. 5279 of *Lecture Notes in Computer Science*. Springer Berlin Heidelberg, 2008, pp. 27–40. 28, 38

[Mic14a] MICROCHIP INC.: Microchip GestIC, 2014. http://www.microchip.com/gestic (accessed 2014-10-24). 15, 79

[Mic14b] MICROSOFT CORP.: Discover the Samsung SUR40, 2014. http://www.microsoft.com/en-us/pixelsense/WhatsNew.aspx (date accessed: 2014/10/12). 26, 33, 38, 97

[MMC09] MISTRY P., MAES P., CHANG L.: Wuw - wear ur world: A wearable gestural interface. In *CHI '09 Extended Abstracts on Human Factors in Computing Systems* (New York, NY, USA, 2009), CHI EA '09, ACM, pp. 4111–4116. xiv, 26, 27, 38

[MMFT07] METZGER C., MEYER J., FLEISCH E., TRÖSTER G.: Weight-sensitive foam to monitor product availability on retail shelves. In *Pervasive Computing*, LaMarca A., Langheinrich M., Truong K., (Eds.), vol. 4480 of *Lecture Notes in Computer Science*. Springer Berlin Heidelberg, 2007, pp. 268–279. 35

[MMOM13] MATIC A., MAXHUNI A., OSMANI V., MAYORA O.: Virtual uniforms: Using sound frequencies for grouping individuals. In *Proceedings of the 2013 ACM Conference on Pervasive and Ubiquitous Computing Adjunct Publication* (New York, NY, USA, 2013), UbiComp '13 Adjunct, ACM, pp. 159–162. 30

[MP02] MORRIS S., PARADISO J.: Shoe-integrated sensor system for wireless gait analysis and real-time feedback. In *Engineering in Medicine and Biology, 2002. 24th Annual Conference and the Annual Fall Meeting of the Biomedical Engineering Society EMBS/BMES Conference, 2002. Proceedings of the Second Joint* (Oct 2002), vol. 3, pp. 2468–2469 vol.3. 35, 38

[MR97] MATSUSHITA N., REKIMOTO J.: Holowall: Designing a finger, hand, body, and object sensitive wall. In *Proceedings of the 10th Annual ACM Symposium on User Interface Software and Technology* (New York, NY, USA, 1997), UIST '97, ACM, pp. 209–210. 26

[MRB10] MCGOOKIN D., ROBERTSON E., BREWSTER S.: Clutching at straws: using tangible interaction to provide non-visual access to graphs. In *CHI '10* (2010), ACM, pp. 1715–1724. 73

[MS04] MICHAHELLES F., SCHIELE B.: Sensing opportunities for physical interaction. In *Physical Interaction 2003 Workshop on Real World User Interfaces* (2004), Mobile HCI 2004. xiv, 7, 37, 38

[MSB11] MATTHIAS STRAKA STEFAN HAUSWIESNER M. R., BISCHOF H.: Skeletal graph based human pose estimation in real-time. In *Proceedings of the British Machine Vision Conference* (2011), BMVA Press, pp. 69.1–69.12. http://dx.doi.org/10.5244/C.25.69. 27

[MSSD06] MAURER U., SMAILAGIC A., SIEWIOREK D., DEISHER M.: Activity recognition and monitoring using multiple sensors on different body positions. In *Wearable and Implantable Body Sensor Networks, 2006. BSN 2006. International Workshop on* (April 2006), pp. 4 pp.–116. 34, 38

[MSW02] MILSTEIN A., SÁNCHEZ J. N., WILLIAMSON E. T.: Robust Global Localization Using Clustered Particle Filtering. *Artificial Intelligence '02* (2002), 581–586. 88

[Mur11] MURATA MANUFACTURING CO, LTD.: Capacitive Coupling Power Transmission Module, 2011. http://www.murata.com/products/article/pdf/ta1291.pdf (accessed 2014-02-25). 69, 72

[Nel05] NELSON M.: Target detection, image analysis, and modeling. In *Electroreception*, Bullock T., Hopkins C., Popper A., Fay R., (Eds.), vol. 21 of *Springer Handbook of Auditory Research*. Springer New York, 2005, pp. 290–317. 9

[NFC14] NFC Forum: Protocol Technical Specification, 2014. http://members.nfc-forum.org/specs/spec_list/ (accessed 2014-02-25). 34, 71

[NLLP04] Ni L., Liu Y., Lau Y., Patil A.: Landmarc: Indoor location sensing using active rfid. *Wireless Networks 10*, 6 (2004), 701–710. 34, 38

[Nor02] Norman D.: *The Design of Everyday Things*. Basic Books, 2002. 8, 123, 124

[NSM*11] Nakatsuma K., Shinoda H., Makino Y., Sato K., Maeno T.: Touch interface on back of the hand. In *ACM SIGGRAPH 2011 Emerging Technologies* (New York, NY, USA, 2011), SIGGRAPH '11, ACM, pp. 19:1–19:1. 33, 38

[OGT*13] Olberding S., Gong N.-W., Tiab J., Paradiso J. A., Steimle J.: A cuttable multi-touch sensor. In *Proceedings of the 26th Annual ACM Symposium on User Interface Software and Technology* (New York, NY, USA, 2013), UIST '13, ACM, pp. 245–254. 21

[OSOI12] Ogata M., Sugiura Y., Osawa H., Imai M.: iring: Intelligent ring using infrared reflection. In *Proceedings of the 25th Annual ACM Symposium on User Interface Software and Technology* (New York, NY, USA, 2012), UIST '12, ACM, pp. 131–136. 33, 38

[OST13] Ono M., Shizuki B., Tanaka J.: Touch & activate: Adding interactivity to existing objects using active acoustic sensing. In *Proceedings of the 26th Annual ACM Symposium on User Interface Software and Technology* (New York, NY, USA, 2013), UIST '13, ACM, pp. 31–40. 135

[PAHR97] Paradiso J., Abler C., Hsiao K.-y., Reynolds M.: The magic carpet: Physical sensing for immersive environments. In *CHI '97 Extended Abstracts on Human Factors in Computing Systems* (New York, NY, USA, 1997), CHI EA '97, ACM, pp. 277–278. 32, 36, 37, 38

[PCB00] Priyantha N. B., Chakraborty A., Balakrishnan H.: The cricket location-support system. In *Proceedings of the 6th Annual International Conference on Mobile Computing and Networking* (New York, NY, USA, 2000), MobiCom '00, ACM, pp. 32–43. 30

[PF01] Paradiso J., Feldmeier M.: A compact, wireless, self-powered pushbutton controller. In *Ubicomp 2001: Ubiquitous Computing*, Abowd G., Brumitt B., Shafer S., (Eds.), vol. 2201 of *Lecture Notes in Computer Science*. Springer Berlin Heidelberg, 2001, pp. 299–304. 36, 38

[PFKP05] Patterson D., Fox D., Kautz H., Philipose M.: Fine-grained activity recognition by aggregating abstract object usage. In *Proceedings of the 9th IEEE International Symposium on Wearable Computers (ISWC '05)* (2005), pp. 44–51. 104

[Pfl95] Pflüger E.: *Archiv für die gesammte Physiologie des Menschen und der Thiere: 1895*. No. Bd. 60. M. Hager, 1895. 9

[PFP*04] Philipose M., Fishkin K. P., Perkowitz M., Patterson D. J., Fox D., Kautz H., Hahnel D.: Inferring activities from interactions with objects. *IEEE Pervasive Computing 3*, 4 (Oct. 2004), 50–57. 104

[PG97] Paradiso J. A., Gershenfeld N.: Musical applications of electric field sensing. *Computer Music Journal 21* (1997), 69–89. 10, 38

[PGGP13] Pu Q., Gupta S., Gollakota S., Patel S.: Whole-home gesture recognition using wireless signals. In *Proceedings of the 19th Annual International Conference on Mobile Computing & Networking* (New York, NY, USA, 2013), MobiCom '13, ACM, pp. 27–38. 32, 37, 38, 135

[PH97] PARADISO J. A., HU E.: Expressive footwear for computer-augmented dance performance. In *Proceedings of the 1st IEEE International Symposium on Wearable Computers* (Washington, DC, USA, 1997), ISWC '97, IEEE Computer Society, pp. 165–. 36, 37, 38

[PH14] PEDERSEN E. W., HORNBÆK K.: Expressive touch: Studying tapping force on tabletops. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems* (New York, NY, USA, 2014), CHI '14, ACM, pp. 421–430. 29

[PKS*06] PARK D. G., KIM J. K., SUNG J. B., HWANG J. H., HYUNG C. H., KANG S. W.: TAP: touch-and-play. In *CHI '06* (2006), ACM, pp. 677–680. 25, 70

[PLCH02] PARADISO J., LEO C. K., CHECKA N., HSIAO K.: Passive acoustic sensing for tracking knocks atop large interactive displays. In *Sensors, 2002. Proceedings of IEEE* (2002), vol. 1, pp. 521–527 vol.1. 29, 36, 37, 38, 40

[PLSR08] POPESCU M., LI Y., SKUBIC M., RANTZ M.: An acoustic fall detector system that uses sound height information to reduce the false alarm rate. In *Engineering in Medicine and Biology Society, 2008. EMBS 2008. 30th Annual International Conference of the IEEE* (Aug 2008), pp. 4628–4631. 29, 38

[Pra14] PRATT S.: Ask the Application Engineer - Capacitance Sensors for Human Interfaces to Electronic Equipment, 2014. http://www.analog.com/TODO (date accessed: 2014/08/25). xiii, 21

[Pre10] PRESS M.: PrimeSense Supplies 3-D-Sensing Technology to 'Project Natal' for Xbox 360. `http://www.primesense.com/en/press-releases/press-releases2010/109-primesense` (accessed 30/09/2014). 27, 28, 38, 125

[PS05] PARADISO J., STARNER T.: Energy scavenging for mobile and wireless electronics. *Pervasive Computing, IEEE* (2005), 18–27. 72

[PSH97] PAVLOVIC V., SHARMA R., HUANG T.: Visual interpretation of hand gestures for human-computer interaction: a review. *IEEE Transactions on Pattern Analysis and Machine Intelligence 19*, 7 (July 1997), 677–695. 55

[PSLS12] POUPYREV I., SCHOESSLER P., LOH J., SATO M.: Botanicus interacticus: interactive plants technology. In *SIGGRAPH Emerging Technologies* (2012), ACM, p. 4. 22

[PSS*11] PFAFFEN S., SOMMER P., STOCKER C., WATTENHOFER R., WELTEN S.: Planipes: Mobile foot pressure analysis. In *Proceedings of the First ACM Workshop on Mobile Systems, Applications, and Services for Healthcare* (New York, NY, USA, 2011), mHealthSys '11, ACM, pp. 2:1–2:6. 35, 38

[PSZ12] PENG C., SHEN G., ZHANG Y.: Beepbeep: A high-accuracy acoustic-based system for ranging and localization using cots devices. *ACM Trans. Embed. Comput. Syst. 11*, 1 (Apr. 2012), 4:1–4:29. 30, 37

[PYGH10] POUPYREV I., YEO Z., GRIFFIN J. D., HUDSON S.: Sensing human activities with resonant tuning. In *CHI '10 Extended Abstracts on Human Factors in Computing Systems* (New York, NY, USA, 2010), CHI EA '10, ACM, pp. 4135–4140. 22

[Rab89] RABINER L. R.: A tutorial on hidden markov models and selected applications in speech recognition. In *Proceedings of the IEEE* (1989), vol. 77, pp. 257–286. 90, 91

[RCP10] RALSTON T., CHARVAT G., PEABODY J.: Real-time through-wall imaging using an ultrawideband multiple-input multiple-output (mimo) phased array radar system. In

Phased Array Systems and Technology (ARRAY), 2010 IEEE International Symposium on (Oct 2010), pp. 551–558. 32, 38

[RDML05] RAVI N., DANDEKAR N., MYSORE P., LITTMAN M.: Activity recognition from accelerometer data. In *Proceedings of the National Conference on Artificial Intelligence* (2005), vol. 20, Menlo Park, CA; Cambridge, MA; London; AAAI Press; MIT Press; 1999, p. 1541. 104

[Rek97] REKIMOTO J.: Pick-and-drop: A direct manipulation technique for multiple computer environments. In *Proceedings of the 10th Annual ACM Symposium on User Interface Software and Technology* (New York, NY, USA, 1997), UIST '97, ACM, pp. 31–39. 136

[RGPK14] RUS S., GROSSE-PUPPENDAHL T., KUIJPER A.: Recognition of Bed Postures using Mutual-Capacitance Sensing. In *Ambient Intelligence* (2014), Lecture Notes in Computer Science, Springer. 57

[RHR10] RIEDENKLAU E., HERMANN T., RITTER H.: Tangible Active Objects and Interactive Sonification as a Scatter Plot Alternative for the Visually Impaired. In *International Conference on Auditory Display* (2010), ACM. 73

[RM01] RANDELL C., MULLER H. L.: Low cost indoor positioning system. In *Proceedings of the 3rd International Conference on Ubiquitous Computing* (London, UK, UK, 2001), UbiComp '01, Springer-Verlag, pp. 42–48. 30, 38

[RMD07] REYNOLDS M., MAZALEK A., DAVENPORT G.: An acoustic position sensing system for large scale interactive displays. In *Sensors, 2007 IEEE* (Oct 2007), pp. 1193–1196. 30, 37, 38, 136

[RMSAR07] ROUGIER C., MEUNIER J., ST-ARNAUD A., ROUSSEAU J.: Fall detection from human shape and motion history using video surveillance. In *Advanced Information Networking and Applications Workshops, 2007, AINAW '07. 21st International Conference on* (May 2007), vol. 2, pp. 875–880. 27, 37, 38

[Rus13] RUS S.: *Recognition of Lying Postures using Capacitive Proximity Sensing*. Master's thesis, Technische Universität Darmstadt, 2013. 57

[SBK*11] SCOTT J., BRUSH A. B., KRUMM J., MEYERS B., HAZAS M., HODGES S., VILLAR N.: Preheat: Controlling home heating using occupancy prediction. In *Proceedings of UbiComp 2011* (2011), ACM. 137

[SBW12] SODHI R., BENKO H., WILSON A.: Lightguide: projected visualizations for hand movement guidance. In *CHI '12* (2012), pp. 179–188. 125

[SCC10] SRINIVASAN R., CHEN C., COOK D.: Activity recognition using actigraph sensor. In *Proceedings of the Fourth Int. Workshop on Knowledge Discovery form Sensor Data (ACM SensorKDD'10), Washington, DC, July* (2010), pp. 25–28. 104

[Sch00] SCHMIDT A.: Implicit human computer interaction through context. *Personal Technologies 4*, 2-3 (2000), 191–199. v, 8, 103, 123

[SCZ*14] SUNDHOLM M., CHENG J., ZHOU B., SETHI A., LUKOWICZ P.: Smart-mat: Recognizing and counting gym exercises with low-cost resistive pressure sensing matrix. In *Proceedings of the 2014 ACM International Joint Conference on Pervasive and Ubiquitous Computing* (New York, NY, USA, 2014), UbiComp '14, ACM, pp. 373–382. xiv, 35, 37, 38

[SFJ*05] SMITH J. R., FISHKIN K. P., JIANG B., MAMISHEV A., PHILIPOSE M., REA A. D., ROY S., SUNDARA-RAJAN K.: Rfid-based techniques for human-activity detection. *Commun. ACM 48*, 9 (Sept. 2005), 39–44. 34, 38, 39

[SGB99] SMITH J. R., GERSHENFELD N., BENTON S. A.: Electric Field Imaging. *Technology* (1999). xiii, 3, 13, 16, 23, 37, 38, 49, 52, 65, 69, 70, 79, 80, 81, 84, 85, 94, 126

[SHS01] SAVVIDES A., HAN C.-C., STRIVASTAVA M. B.: Dynamic fine-grained localization in ad-hoc networks of sensors. In *Proceedings of the 7th Annual International Conference on Mobile Computing and Networking* (New York, NY, USA, 2001), MobiCom '01, ACM, pp. 166–179. 30, 38

[SHVLS08a] STIKIC M., HUYNH T., VAN LAERHOVEN K., SCHIELE B.: ADL Recognition Based on the Combination of RFID and Accelerometer Sensing. In *Pervasive Health 2008* (January 2008), IEEE, p. 258–263. 76

[SHVLS08b] STIKIC M., HUYNH T., VAN LAERHOVEN K., SCHIELE B.: ADL Recognition Based on the Combination of RFID and Accelerometer Sensing. In *Proceedings of the 2nd International Conference on Pervasive Computing Technologies for Healthcare (Pervasive Health 2008)* (Tampere, Finland, January 2008), IEEE Xplore, IEEE Xplore, p. 258–263. 104

[SL08] STEINHAGE A., LAUTERBACH C.: Sensfloor (r): Ein aal sensorsystem für sicherheit, homecare und komfort. In *AAL Kongress '08* (2008). (in German). vi, 3, 56

[Smi96] SMITH J. R.: Field mice: Extracting hand geometry from electric field measurements. *IBM Syst. J. 35*, 3-4 (Sep 1996), 587–608. xiii, 11, 12, 14, 15, 16, 19, 20, 23, 24, 79, 81, 85

[SN11] SAAD S., NAKAD Z.: A standalone rfid indoor positioning system using passive tags. *Industrial Electronics, IEEE Transactions on 58*, 5 (May 2011), 1961–1970. 34

[SPGI13] SODHI R., POUPYREV I., GLISSON M., ISRAR A.: Aireal: Interactive tactile experiences in free air. *ACM Trans. Graph. 32*, 4 (July 2013), 134:1–134:10. 125, 136

[SPH12] SATO M., POUPYREV I., HARRISON C.: Touché: enhancing touch interaction on humans, screens, liquids, and everyday objects. In *CHI '12* (2012), pp. 483–492. 18, 22, 38, 70, 79

[SRO*08] STIEFMEIER T., ROGGEN D., OGRIS G., LUKOWICZ P., TRÖSTER G.: Wearable activity tracking in car manu- facturing. *IEEE Pervasive Computing 7* (2008), 42–50. 104

[SSvL12] SCHMIDT A., SHIRAZI A., VAN LAERHOVEN K.: Are you in bed with technology? *Pervasive Computing, IEEE 11*, 4 (Oct 2012), 4–7. 76

[Sta14] STALLER E.: Spirogyrate, 2014. http://ericstaller.com/commissioned%20work/SPIROGYRATE/ (date accessed: 2014/10/29). xv, 56

[STM14] STMICROELECTRONICS INC.: STRFNFCA Near field communication transceiver, 2014. http://www.st.com/st-web-ui/static/active/en/resource/technical/document/datasheet/DM00052610.pdf (accessed 2014-06-17). 71

[STS*13] SOUSA M., TECHMER A., STEINHAGE A., LAUTERBACH C., LUKOWICZ P.: Human tracking and identification using a sensitive floor and wearable accelerometers. In *PerCom '13* (2013), pp. 166–171. 70

[SVL01] SCHMIDT A., VAN LAERHOVEN K.: How to build smart appliances? *Personal Communications, IEEE 8*, 4 (Aug 2001), 66–71. 103

[SYP*08] SAMPLE A., YEAGER D., POWLEDGE P., MAMISHEV A., SMITH J.: Design of an rfid-based battery-free programmable sensing platform. *Instrumentation and Measurement, IEEE Transactions on 57*, 11 (Nov 2008), 2608–2615. 34, 38

[SZH12] SAVAGE V., ZHANG X., HARTMANN B.: Midas: Fabricating Custom Capacitive Touch Sensors to Prototype Interactive Objects. In *UIST '12* (2012). 21, 44, 70

[TDDM09] TARZIA S. P., DICK R. P., DINDA P. A., MEMIK G.: Sonar-based measurement of user presence and attention. In *Proceedings of the 11th International Conference on Ubiquitous Computing* (New York, NY, USA, 2009), Ubicomp '09, ACM, pp. 89–92. 30, 38

[Tex14a] TEXAS INSTRUMENTS: LMC555 Datasheet, 2014. http://www.ti.com/lit/ds/symlink/lmc555.pdf (date accessed: 2014/08/30). 48

[Tex14b] TEXAS INSTRUMENTS INC.: Application Report SLOA184 - NFC and RFID Reader Ultra-Low-Power Card Presence Detection Using MSP430 and TRF79xxA, 2014. http://www.ti.com/lit/an/sloa184/sloa184.pdf (accessed 2014-02-25). 71

[Tex14c] TEXAS INSTRUMENTS INC.: TRF7970A Multiprotocol Fully Integrated 13.56-MHz RFID and Near Field Communication (NFC) Transceiver IC, 2014. http://www.ti.com/lit/ds/slos743k/slos743k.pdf (accessed 2014-02-25). 71

[TITO11] TAKEMURA K., ITO A., TAKAMATSU J., OGASAWARA T.: Active bone-conducted sound sensing for wearable interfaces. In *Proceedings of the 24th Annual ACM Symposium Adjunct on User Interface Software and Technology* (New York, NY, USA, 2011), UIST '11 Adjunct, ACM, pp. 53–54. 29, 38

[TR00] TURK M., ROBERTSON G.: Perceptual user interfaces (introduction). *Commun. ACM 43*, 3 (Mar. 2000), 32–34. 8

[TTW*05] TAO J., TURJO M., WONG M.-F., WANG M., TAN Y.-P.: Fall incidents detection for intelligent video surveillance. In *Information, Communications and Signal Processing, 2005 Fifth International Conference on* (2005), pp. 1590–1594. 26, 27, 38

[Uni11] UNIVERSITY OF WAIKATO: Weka machine learning project, May 2011. 115

[VDP03] VERMAAK J., DOUCET A., PATRICK P.: Maintaining Multi-Modality through Mixture Tracking. In *ICCV '03* (2003), pp. 1110–1118. 88

[VG13] VU T., GRUTESER M.: Personal Touch-Identification Tokens. *IEEE Pervasive Computing 12*, 2 (2013), 10–13. xiv, 4, 25, 38, 70, 136

[VM92] VARNISH J., MCCONNELL R.: Driven shielding capacitive proximity sensor, Dec. 7 1992. CA Patent App. CA 2,070,309. 14, 16

[VMV09] VALTONEN M., MÄENTAUSTA J., VANHALA J.: Tiletrack : Capacitive human tracking using floor tiles. In *PerCom '09* (2009), pp. 121–133. vi, 3, 24, 56, 57, 70

[VSH11] VILLAR N., SCOTT J., HODGES S.: Prototyping with microsoft .net gadgeteer. In *Proceedings of the Fifth International Conference on Tangible, Embedded, and Embodied Interaction* (2011), ACM, p. 377–380. 43

[WB10] WILSON A. D., BENKO H.: Combining multiple depth cameras and projectors for interactions on, above and between surfaces. In *Proceedings of the 23Nd Annual ACM Symposium on User Interface Software and Technology* (New York, NY, USA, 2010), UIST '10, ACM, pp. 273–282. 136

[WB11] WIMMER R., BAUDISCH P.: Modular and deformable touch-sensitive surfaces based on time domain reflectometry. In *Proceedings of the 24th Annual ACM Symposium on User Interface Software and Technology* (New York, NY, USA, 2011), UIST '11, ACM, pp. 517–526. http://doi.acm.org/10.1145/2047196.2047264. xiv, 31, 37, 38, 70

[WBRF13] WEICHERT F., BACHMANN D., RUDAK B., FISSELER D.: Analysis of the accuracy and robustness of the leap motion controller. *Sensors 13*, 5 (2013), 6380–6393. 26, 38, 97

[Web10] WEBSTER J.: *Medical instrumentation: application and design*. Wiley, 2010. 18

[Wei99] WEISER M.: The computer for the 21st century. *SIGMOBILE Mob. Comput. Commun. Rev. 3*, 3 (July 1999), 3–11. v, 1, 7, 43, 103, 133

[WHFaG92] WANT R., HOPPER A., FALCÃO V., GIBBONS J.: The active badge location system. *ACM Trans. Inf. Syst. 10*, 1 (Jan. 1992), 91–102. 33, 38

[WHKS06] WIMMER R., HOLLEIS P., KRANZ M., SCHMIDT A.: Thracker - using capacitive sensing for gesture recognition. In *Distributed Computing Systems Workshops, 2006. ICDCS Workshops 2006. 26th IEEE International Conference on* (July 2006), pp. 64–64. xiv, 23, 24

[Wil04] WILSON A. D.: Touchlight: An imaging touch screen and display for gesture-based interaction. In *Proceedings of the 6th International Conference on Multimodal Interfaces* (New York, NY, USA, 2004), ICMI '04, ACM, pp. 69–76. 26, 38

[Wil05] WILSON A. D.: Playanywhere: A compact interactive tabletop projection-vision system. In *Proceedings of the 18th Annual ACM Symposium on User Interface Software and Technology* (New York, NY, USA, 2005), UIST '05, ACM, pp. 83–92. 26, 38

[Wim11] WIMMER R.: Grasp sensing for human-computer interaction. In *Proceedings of the Fifth International Conference on Tangible, Embedded, and Embodied Interaction* (New York, NY, USA, 2011), TEI '11, ACM, pp. 221–228. 22

[WKBS07a] WIMMER R., KRANZ M., BORING S., SCHMIDT A.: A Capacitive Sensing Toolkit for Pervasive Activity Detection and Recognition. *Fifth Annual IEEE International Conference on Pervasive Computing and Communications (PerCom'07)* (2007), 171–180. xiv, xv, 3, 23, 24, 37, 38, 44, 45, 51, 52, 54, 55, 57, 70, 94, 135

[WKBS07b] WIMMER R., KRANZ M., BORING S., SCHMIDT A.: Captable and capshelf - unobtrusive activity recognition using networked capacitive sensors. In *INSS '07* (2007), pp. 85–88. vi, xiv, 3, 16, 24, 25, 47

[WLTS06] WARD J., LUKOWICZ P., TROSTER G., STARNER T.: Activity recognition of assembly tasks using body-worn microphones and accelerometers. *Pattern Analysis and Machine Intelligence, IEEE Transactions on 28*, 10 (Oct. 2006), 1553–1567. 104

[WM10] WYSS T., MADER U.: Recognition of Military-Specific Physical Activities With Body-Fixed Sensors. *Military Medicine 175*, 11 (2010), 858–864. 104

[WO05] WILSON A., OLIVER N.: Multimodal Sensing for Explicit and Implicit Interaction. 7

[WT63] WATANABE A., TAKEDA K.: The change of discharge frequency by a.c. stimulus in a weak electric fish. *Journal of Experimental Biology 40*, 1 (1963), 57–66. 9

[XBY*11] XIAO X., BERNSTEIN M. S., YAO L., LAKATOS D., GUST L., ACQUAH K., ISHII H.: Pingpong++: Community customization in games and entertainment. In *Proceedings of the 8th International Conference on Advances in Computer Entertainment Technology* (New York, NY, USA, 2011), ACE '11, ACM, pp. 24:1–24:6. 29, 38

[YCL*11] YU N.-H., CHAN L.-W., LAU S. Y., TSAI S.-S., HSIAO I., TSAI D.-J., HSIAO F.-I., CHENG L.-P., CHEN M., HUANG P., ET AL.: TUIC: enabling tangible interaction on capacitive multi-touch displays. In *CHI '11* (2011), ACM, pp. 2995–3004. 4, 38, 70, 136

[YN12] YAMAKAWA S., NOJIMA T.: A proposal for a mmg-based hand gesture recognition method. In *Adjunct Proceedings of the 25th Annual ACM Symposium on User Interface Software and Technology* (New York, NY, USA, 2012), UIST Adjunct Proceedings '12, ACM, pp. 89–90. 29, 38

[YNR13] YEO K. P., NANAYAKKARA S., RANSIRI S.: Stickear: Making everyday objects respond to sound. In *Proceedings of the 26th Annual ACM Symposium on User Interface Software and Technology* (New York, NY, USA, 2013), UIST '13, ACM, pp. 221–226. 29, 38

[YSC*11] YANG J., SIDHOM S., CHANDRASEKARAN G., VU T., LIU H., CECAN N., CHEN Y., GRUTESER M., MARTIN R. P.: Detecting driver phone use leveraging car speakers. In *Proceedings of the 17th Annual International Conference on Mobile Computing and Networking* (New York, NY, USA, 2011), MobiCom '11, ACM, pp. 97–108. 30

[YT12] YATANI K., TRUONG K. N.: Bodyscope: A wearable acoustic sensor for activity recognition. In *Proceedings of the 2012 ACM Conference on Ubiquitous Computing* (New York, NY, USA, 2012), UbiComp '12, ACM, pp. 341–350. 29, 38

[ZCA*14] ZHAO C., CHEN K.-Y., AUMI M. T. I., PATEL S., REYNOLDS M. S.: Sideswipe: Detecting in-air gestures around mobile devices using actual gsm signal. In *Proceedings of the 27th Annual ACM Symposium on User Interface Software and Technology* (New York, NY, USA, 2014), UIST '14, ACM, pp. 527–534. 32, 38

[Zim96] ZIMMERMAN T. G.: Personal Area Networks (PAN): Near-Field Intra-Body Communication. *IBM Systems Journal 35*, 3&4 (1996), 609–617. xiii, 4, 16, 17, 18, 25, 38, 59, 61, 62, 63, 66, 70

[ZMB*14] ZEISS S., MARINC A., BRAUN A., GROSSE-PUPPENDAHL T., BECK S.: A gesture-based door control using capacitive sensors. In *Distributed, Ambient, and Pervasive Interactions*, Streitz N., Markopoulos P., (Eds.), vol. 8530 of *Lecture Notes in Computer Science*. Springer International Publishing, 2014, pp. 207–216. 57

[ZS13] ZHAO Y., SMITH J.: A battery-free rfid-based indoor acoustic localization platform. In *RFID (RFID), 2013 IEEE International Conference on* (April 2013), pp. 110–117. 34

[ZSP*95] ZIMMERMAN T. G., SMITH J. R., PARADISO J. A., ALLPORT D., GERSHENFELD N.: Applying electric field sensing to human-computer interfaces. In *CHI '95* (1995), pp. 280–287. 15, 16, 23, 24, 51, 55, 61

"It may be the warriors who get the glory. But it's the engineers who build societies."

— B'Elanna Torres, stardate 54337.5