Diss. ETH Nr. 17850

Scenario-Driven Prototyping for Ubiquitous Computing

A Novel Method for the Assessment of Technological Challenges and Societal Implications

A dissertation submitted to the
SWISS FEDERAL INSTITUTE OF TECHNOLOGY ZURICH

for the degree of
Doctor of Sciences

presented by
Vlad-Constantin Coroamă
Diplom-Informatiker, Technical University of Darmstadt
born February 2, 1974
citizen of Romania

accepted on the recommendation of
Prof. Dr. Friedemann Mattern, examiner
Prof. Dr. Lorenz Hilty, co-examiner
Prof. Dr. Thomas Gross, co-examiner

2008

To Denise
for all the love she brings into this world

and remembering Nana
who knew it all from the beginning

Abstract

With the recent rapid advances in sensing technologies, storage capacity, wireless communication, and processing power, the decade-old vision of smart environments and smart everyday items is becoming increasingly feasible. Given its vast potential, the vision is slowly moving – under different names such as "ubiquitous computing," "smart environment," "ambient intelligence," or "internet of things" – from small, isolated research prototypes into large national and international research collaborations, public policy hearings, and the consciousness of the public at large. However, while the technologies comprising ubiquitous computing continue to advance steadily, a gap is widening between the visions depicting their possible usages, and the reality of existing projects and prototypes. In contrast to the prototypes that usually focus on individual technological aspects, ubiquitous computing visions assume a holistic paradigm, with smart environments governing most aspects of our everyday lives on a 24/7 basis, leading not only to complex technological challenges, but also to broad, far-reaching, and difficult to anticipate societal consequences. It is thus hardly surprising that, as reports about ubiquitous computing research have entered into the public consciousness, concerns over the potential risks have increased.

Traditional methods of technology assessment, and in particular scenario analysis, however, are often overwhelmed by the technological broadness and the high dynamism of the field, resulting in technologically naive scenarios that overestimate future technical developments or underestimate the interplay between different technologies. Consequently, few relevant societal issues are raised, and hardly any novel technological insights and solutions can be drawn from such exercises. Developers of ubiquitous computing technologies, on the other hand, have traditionally focused mostly on purely technological challenges in their research projects and prototypes, e.g., smart office environments, novel location technologies, or efficient service discovery protocols. The narrow scope of such prototypes typically ignores questions of societal compatibility and economic feasibility, thus offering no help in understanding the broad societal opportunities and risks of the novel technologies either. Furthermore, since the projects often ignore societal and economic constraints, many of the technological answers they provide could prove to be irrelevant, for such applications might never exist in the future. If ubiquitous computing prototypes were to focus more on recognizing relevant areas – i.e., applications with a solid business model, strong societal acceptance, and few societal risks – they would not only provide the needed help in understanding the societal challenges, but would also approach more relevant technological issues.

The aim of this thesis is to define a novel method for identifying the relevant challenges posed by ubiquitous computing technologies – both technical and societal challenges. The method, “scenario-driven prototyping,” proposes in a four-step process the tight coupling of multidisciplinary scenario analysis with technological prototype development. The hypothesis is that closer coordination of scenario forecasting and prototype development can be combined in an integrated, coordinated process, thus helping to narrow the above-mentioned gap between technological feasibility and societal concern. This tighter integration will allow for relevant prototypes that advance technological progress through more realistic challenges. At the same time, the better grounding of scenarios in technological realities will help alleviate needless fears from misrepresented technical abilities, and focus instead on the probable opportunities and risks posed by ubiquitous computing.

To this respect, the contributions of the thesis are threefold: First, the scenario-driven prototyping method itself. Second, the technological and societal insights gained from an assistive technology project devoted to the blind and visually impaired, which was developed according to the proposed method. Finally, the insights from a project for the individual and behavior-dependent accounting of traffic costs, also developed according to that method. Through the multitude of interdependent technological, societal, and economic insights generated, many of which would presumably not have been possible with formerly existing methods, both projects support our claim that scenario-driven prototyping is a valuable tool to address relevant technological and societal implications of ubiquitous computing.

Zusammenfassung

Mit den raschen Fortschritten der letzten Jahre scheint nun die über 15 Jahre alte Vision smarter Umgebungen und smarter Alltagsgegenstände realistischer denn je. Diese Vision, ursprünglich von Mark Weiser unter dem Namen „Ubiquitous Computing“ geprägt, findet sich nun – unter verschiedenen weiteren Namen, etwa „smarte Umgebung“, „ambient intelligence“ oder „Internet der Dinge“ – in verschiedenen Forschungsprogrammen und gelangt auch immer mehr in das Bewusstsein eines breiten Publikums. Während die Technologien des Ubiquitous Computing kontinuierlich fortschreiten, öffnet sich allerdings mehr und mehr eine Schere zwischen den Visionen, die mögliche Anwendungen des Ubiquitous Computing beschreiben, und der Realität existierender Projekte und Prototypen. Während nämlich die zahlreichen Projekte meist einzelne Einsatzmöglichkeiten mit beschränkter Reichweite hervorheben, wird Ubiquitous Computing in den Visionen als ein ganzheitlich ausgerichtetes Paradigma verstanden. Sollten sich derartige Visionen bewahrheiten, so würden die vom Ubiquitous Computing ermöglichten smarten Umgebungen letztlich die verschiedensten Aspekte unseres Alltagslebens rund um die Uhr bestimmen. Um die damit verbundene Komplexität zu meistern, müssten nicht nur die technologischen Herausforderungen berücksichtigt werden, sondern auch auf die zahlreichen auftauchenden gesellschaftlichen Fragen und Vorbehalte Antworten gegeben werden können.

Das klassische Instrument der Technikfolgenabschätzung – Zunkunftsszenarien – scheint jedoch für Ubiquitous Computing in weiten Teilen zu versagen. Verantwortlich dafür ist die hohe Komplexität der Aufgabe, welche in der Breite des Feldes, der Fülle an Lebensbereichen, die von Ubiquitous-Computing-Technologien beeinflusst werden könnten, sowie in der hohen Dynamik technologischer Innovationen in diesem Bereich begründet liegt. Oft beruhen Szenarien nur auf der Analyse wenig relevanter Anwendungsfelder oder auf einem oberflächlichen Verständnis technologischer Möglichkeiten, so dass die von ihnen gelieferten Ergebnisse zum Teil irreführend sind. Die Entwickler von Ubiquitous-

Computing-Technologien, andererseits, könnten an sich die benötigten realistischen Einschätzungen der aktuellen Möglichkeiten sowie der künftigen technologischen Entwicklungsrichtungen liefern. Die meisten Forschungsprojekte und -Prototypen, obwohl zahlreiche technische Herausforderungen angehend, haben jedoch auch nicht wesentlich zur Bewältigung der gesellschaftlichen Herausforderungen beitragen können. Da sie oft von einem engen Horizont charakterisiert sind und typischerweise soziale Verträglichkeit und ökonomischen Nutzen ignorieren, sind sie nicht gut dazu geeignet, die umfassende Frage gesellschaftlicher Chancen und Risiken zu beantworten. Darüber hinaus könnten sich sogar einige der in solchen Projekten gelieferten technischen Beiträge letztlich als irrelevant in der Anwendungspraxis erweisen. Durch die fehlende gesellschaftliche und ökonomische Verankerung der Prototypen werden viele dieser Anwendungen nie existieren, so dass es ebenfalls keiner Lösung für die von ihnen aufgeworfenen Probleme bedarf. Falls jedoch Ubiquitous-Computing-Prototypen mehr auf relevante Bereiche achten würden – Anwendungen mit einem soliden ökonomischen Modell und einem gesellschaftlichem Nutzen – könnten sie nicht nur besser der Untersuchung gesellschaftlicher Herausforderungen dienen, sondern auch die relevanten technischen Aspekte evozieren und untersuchen.

Das Ziel der vorliegenden Arbeit ist es daher, eine neuartige Methode zu präsentieren, die der Identifizierung relevanter Herausforderungen – technologischer wie auch gesellschaftlicher Art – des Ubiquitous Computing dient. Die Methode der „Szenarien-getriebener Prototypentwicklung“ („scenario-driven prototyping“) verfolgt in einem vierstufigen Prozess die enge Verknüpfung technologischer Prototypentwicklung und interdisziplinärer Szenarioanalyse. Die Hypothese dabei ist, dass durch diese enge Verzahnung von Szenarien- und Prototypentwicklung die oben erwähnte Diskrepanz zwischen technologischer Machbarkeit und gesellschaftlichen Herausforderungen von beiden Seiten her entschärft werden kann. Die Entwicklung von Projekten nach einer eingehenden Szenarioanalyse sollte durch die resultierenden Prototypen sowohl relevantere technische Fragestellungen aufwerfen und so das Gebiet des Ubiquitous Computing vorantreiben, wie auch den Szenarien selbst eine nützliche technische Rückkopplung bieten. Dadurch wird das Risiko minimiert, dass die Analyse möglicher gesellschaftlicher Auswirkungen von Ubiquitous Computing auf einer dünnen, teils unrealistischen technologischen Basis fundiert und so unnötige Bedenken hervorruft. Stattdessen kann die Szenarienanalyse auf die Fülle realistischer Chancen

und Risiken fokussiert werden.

Die vorliegende Arbeit liefert drei wesentliche Beiträge: Erstens, die Methode der Szenarien-getriebenen Prototypentwicklung selbst. Zweitens, die Resultate eines ersten Projekts, das nach der Methode der Szenarien-getriebenen Prototypentwicklung realisiert wurde, die „gesprächige Umgebung“ („Chatty Environment“). Die gesprächige Umgebung ist ein Projekt zur Unterstützung Blinder und visuell Behinderter im Alltag, welches die erwartete zunehmende Ausbreitung smarter Gegenstände benutzt, um Sehbehinderten Informationen zu liefern, zu denen sie heute keinen Zugang haben, mit dem Ziel, ihnen dadurch ein unabhängigeres Leben zu ermöglichen. Drittens, die Resultate eines weiteren Projekts, das nach derselben Methode durchgeführt wurde, der „smarte Fahrtenschreiber“ („Smart Tachograph“). Der smarte Fahrtenschreiber ist eine Ubiquitous-Computing-Infrastruktur zur feingranularer und verursachergerechten Abrechnung verschiedenster Arten von Verkehrskosten, die für unterschiedlichste verkehrspolitische Ziele eingesetzt werden kann.

Durch die Fülle der teilweise zusammenhängenden technologischen, gesellschaftlichen und ökonomischen Erkentnisse, die in den zwei Projekten generiert wurden, von denen einige mit bereits existierenden Methoden vermutlich nicht hätten erzielt werden können, sehen wir uns in der Annahme bestätigt, dass die Methode der Szenarien-getriebenen Prototypentwicklung ein wertvolles Werkzeug darstellt, um die komplexen technologischen Bedingungen und gesellschaftlichen Folgen des Ubiquitous Computing besser zu verstehen.

Acknowledgements

> *If you want them to fight each other, throw them corn. But if you wish them to be brothers, have them build a tower together.*
> Antoine de Saint-Exupery – La Citadelle

My advisor, Prof. Friedemann Mattern, taught me all I know about scientific thinking and writing. We often had distinct opinions, just for me to later realize that he was right. Thank you for it all, Friedemann. And for always keeping your belief in people.

By writing together with Marc Langheinrich, I could steal some of his amazing art of writing – or at least I hope so. A further great source of scientific inspiration and support, especially in the beginning, has been my old friend Andrei Popovici.

I am further grateful to my two co-advisors: Discussing about the thesis for several months with Prof. Lorenz Hilty has been pure intellectual joy, instead of a burden – as they say it should be with one's co-advisors. As for Prof. Thomas Gross, he generously jumped onto the boat on short notice and provided invaluable insights in a short time.

Undoubtedly also due to Friedemann, I was fortunate to be part of one of the most pleasant research groups ever. Special warmth comes up when remembering the first-hour 'distributed simulations' group: Jürgen 'JJ' Bohn, Harald 'Harry' Vogt, Oliver 'Uolli' Kasten, Kay Römer, and Marc Langheinrich. With JJ, I not only shared office for many years, but also a deep friendship; Kay has often proven the generosity below his quiet character; Harry and Uolli, each in its own and special way, have been a large enrichment; and Marc has been the catalyst for it all. Not officially part of our group, Etzard Stolte was a de facto member – be it for the 'Doppelkopf' nights at his place only.

When the PhD time is long enough, one also gets to know many new colleagues. Thanks to Matthias and Mila Ringwald for the open door of their home; to Thomas Schoch for his large heart below the rough surface; to Christian Flörkemeier for the advices, backcountry skiing and sailing; to Robert Adelmann for his readiness to help; to Philipp Bolliger and Benedikt Ostermeier for continuing the D47.2 spirit; to Christian Frank for being so Romanian; to Silvia Santini for being so Italian; to Christof Roduner for being so non-Swiss; and to Steve Hinske for being himself. Christof, Christian Frank, and above all Steve meant, as JJ earlier, friendship and support in difficult times. Special thanks to all of you.

From the students I have worked with, Felix Röthenbacher, Christoph Plüss, and Michael Duller have been any advisor's dream: dedicated, full of ideas, hard working, and, above all, fun to work with. Felix and Christoph not only developed most parts of the two prototypes presented in this work, but meant much more.

From the ETH, I cannot fail to mention Heidi Gülgün and Hans Dubach, who carefully shielded us from the bureaucracy, and Denise Spicher, who has been an angel in a rather critical moment.

It has been a pleasure to work in the Ladenburg collegium, and with all representatives of the Gottlieb Daimler- and Karl Benz-foundation: Dr. Diethard Schade, refined intellectual and continuous source of inspiration; Dr. Jörg Klein, generous and always focusing on the scientific essence; and Susanne Hallenberger, who, in her always interested and caring way, has grown to become very close.

The numerous Ladenburg meetings have been so pleasant also because of the participants: Michael Kreutzer, Adolf Hohl, Moritz Strasser, Rotraud Gitter, Jörg Hähner, Tobias Zimmer, and especially Carsten Magerkurth, Matthias Handy and Jürgen Müller.

Further thanks belong to my former employers at Helbling AG, Andreas Suter and Tom Russi, who understood to give me the time needed to finish the thesis.

Thanks to my North American friends, Amanda Williams and Elaine May Huang for the English, and for making me always look forward to conferences. To my friends in Brazil, Diego Sanchez Gallo, Charles Christian Miers, Marcos Tork, Thomas Püttiker, Jéssica Martins Camargo, Ana Paula Aprígio Assis, Daniel 'Brutus' Daminelli, and Tiago 'Pexe' Hermenegildo, for the talks, rides, drinks, smokes, hikes, trips, and homes that I could share with them in Paradise. And to Emiko and Noribumi Shirai for making me feel home far from home.

Thanks to Rodica Ciobanu for being my eternal moral compass. To Margot Constantin for her warmth. To Anja Kollmar for the countless tango nights. And, of course, to Amalia Hellwig, for all that she is, has done, and continues doing.

To Denise Uyen-Thi Le, for the quiet, beautiful, innocent, and happy years we spent together. For all her care. And, above all, for showing me love.

Thanks to my father Radu Coroamă for the care, the stories, the love for dancing, the world maps in my childhood bedroom, and for Jules Verne. To my mother Olimpia Coroamă for the strength, the freedom, her never-ending belief, love and support, and for Karl May. To both for raising a free spirit in unlikely environments. Thanks to my grandparents: To Sextil Coroamă for the kindness, the childhood trips, and the name. To Sorana 'Nana' Coroamă-Stanca for her love for family, beauty, and refinement, and for the long-gone world that she so gracefully represented. To Gheorghe 'Tuţu' Banu, for his character and for the other long-gone world that he so proudly still represents. And for the bicycle. To Zamfira 'Miţa' Banu for her kindness, healthy morals, and deep wisdom.

And to Leila Teruko Shirai. For the songs, the paintings, and the dreams. For the love and healing seasons. For letting me come from far, take her out of the water, and then share Paradise with me. For Caraguá and Picinguaba. For inspiring and supporting me to follow my true purpose. For showing me infinity.

Contents

1 Introduction

With the recent rapid advances in sensing technologies, storage capacity, wireless communication, and processing power, visions of smart environments now seem feasible, where objects are able to “feel” (sense), “think” (process), “remember” (store), and “talk” (communicate) to each other and to humans. Many societal opportunities and risks come along with this vision of “ubiquitous computing.” The early assessment of possible societal implications of the new technologies is thus of high importance for researchers, politicians, and society. While assessing future implications of technology is inherently uncertain and error-prone, doing so has proven to be particularly difficult for ubiquitous computing, due to its broad scope, which encompasses a large amount of different technologies and virtually any application domain. This thesis presents a novel approach for the assessment of societal consequences of ubiquitous computing, an approach rooted in a three-year interdisciplinary research process. The approach, called scenario-driven prototyping, tightly brings together the technical prototype development with the scenario-based studies, providing the advantages of both approaches plus some specific advantages, which can only emerge from the synergy of the two.

In this introductory chapter, we introduce the field of ubiquitous computing, motivate the importance of analyzing its possible far-reaching societal consequences, and explain where the specific difficulties lie. We outline the main contributions of our thesis and conclude with an overview of the remaining chapters.

1.1 Motivation

One of the first to foresee the development towards computing power that spreads into everyday things, has been Mark Weiser, a researcher from the Xerox Palo Alto research center.[1] His 1991 visionary article “The Computer for the 21st Century” starts with the statement:

---

[1] See www.parc.xerox.com.

"The most profound technologies are those that disappear. They weave themselves into the fabric of everyday life until they are indistinguishable from it" [163], a development that according to Weiser will be undergone by computers over the next few decades. In the same article, he also coined the term *Ubiquitous Computing*, used to describe his vision of omnipresent computers unobtrusively helping people in many aspects of their everyday lives, without even being noticed.

Since Weiser's article, however, more than 15 years have passed, and, although a tremendous amount of progress towards his vision has been realized on technological side (most notably the availability of ubiquitous information through the World Wide Web and the ubiquity of electronic communication devices in form of mobile phones), "many aspects of Mark Weiser's vision of ubiquitous computing appear as futuristic today as they appeared in 1991," as Davies and Gellersen noted [39]. According to the authors, the main reason is that beyond the technical problems commonly encountered by any prototype, there are various profound technological challenges and societal implications faced by ubiquitous computing applications that extend beyond the prototype stage, challenges that have rarely been addressed and so far remain largely unanswered. Many of the prototypes developed within the field have a gadget-like character, and are thus ill-suited to address further-reaching challenges and implications.

Specifically, based on the analysis of Davies and Gellersen, we identified that most ubiquitous computing prototypes have failed to provide for one or both of the following:

- *A solid economic and societal foundation.* Being often developed as technological showcases, they are not concerned with the economic viability of the application depicted or its societal sustainability. Not doing so, however, often renders the prototype unrealistic and thus irrelevant for the study of societal consequences.
- *The choice of a societally relevant application domain*, which would allow the analysis of complex consequences. Often having a narrow scope, addressing only the world directly known to their developers, ubiquitous computing prototypes are thus typically concerned with small convenience improvements in the office life of the typical white-collar worker. Obviously, this will not allow for a far-reaching, complex societal analysis.

However, as long as researchers are unable to provide satisfying answers for the societal opportunities and possible pitfalls of the new technologies, a paradoxical phenomenon occurs: While many of the original promises of ubiquitous computing have not become reality despite the tremendous technological progress, public concerns about the societal impact of these technologies is steadily increasing.

On the other hand, there exist several studies about the impact of ubiquitous computing technologies done by social scientists. They are usually scenario-based, i.e., a fictional scenario depicts a future ubiquitous computing application in a specific application domain, which is then analyzed from the particular perspective of the participating social scientists. Psychological, sociological, economic, or privacy protection implications can for example be analyzed.

However, this approach encounters following types of problems:

- *Being based on irrelevant or unrealistic prototypes.* Since the social scientist have to study a specific application, they usually resort to one of the prototypical applications developed by engineers. These, however, as argued above, are often irrelevant, or economically and societally unrealistic, or both. In the first case, the analysis will not allow for complex conclusions, in the latter, it can easily produce misleading conclusions, since it takes as granted an application class which will most likely not exist.

- *Naive technological understanding.* To avoid such mistakes as above, social scientist sometimes extrapolate from the prototypes done by engineers to imagine other, more relevant applications, which would use the same technologies that had been used for the rather irrelevant prototypes. However, this process is exposed to the limitations in understanding the possibilities of technology by the social scientists. Even small misunderstandings can lead to large misconceptions about the power of the technology, as for example the confusion between passive and active RFID shows. Passive tags are light-weight, can be hidden in almost any objects, but can only be read from a distance of up to a few meters with bulky, difficult-to-hide antennas. Active tags are large, have a battery that needs to be recharged regularly, but can communicate to distances up to one kilometer. If the two get mixed, resulting in a (technologically impossible) light-weight, hard to detect tag, which can be read from kilometers, misleading and thus dangerous Orwellian scenarios quickly arise.

- *Ignoring pragmatic technologic limitations.* Even when neither of the two mistakes above happens, and the scenario depicts a relevant application with (as far as can be foreseen) realistic technologies, the absence of a prototype still has severe drawbacks. As any engineer will have a story or two to tell, only when developing a prototype as close to reality as possible, the various technological problems will show themselves. Even if they are not of fundamental nature (as the RFID tag misconception addressed above), but more pragmatic limitations, they can still mean several years delay for the availability of an application, or costs so high to overcome the limitation, that the application becomes unrealistic.

For the ubiquitous computing assessment, it is thus necessary to identify the applications that are technologically realistic, of likely future importance, with a strong business model behind for economic realism, and societally sustainable. Finding precisely these applications is a non-trivial task for any technology domain. For ubiquitous computing, however, the task is particularly difficult because of its exceptionally broad scope. It is composed of a large set of different, quickly evolving individual technologies, which can find applications in virtually any domain, having finally the potential to influence every aspect of our lives.

1.2 Contributions

To overcome the drawbacks of technological gadget development and technology-detached scenario analysis, we propose in this thesis a close interdisciplinary coupling of the two, inside a method that we call scenario-driven prototyping. This novel paradigm distillates and structures the experiences gathered and the actual developments accomplished while managing the three-year interdisciplinary project "Living in a Smart Environment – Implications of Ubiquitous Computing."[2] Our thesis also encompasses the presentation of two different technical projects in which we have analyzed and presented possible consequences of ubiquitous computing technology by following the scenario-driven prototyping method: The Smart Tachograph, a system for the precise and individual measurement and accounting of traffic costs, and

---

[2] See www.smart-environment.de.

the Chatty Environment, a system providing the blind and visually impaired more everyday independence.

The major contributions of this thesis, which have also been published in conferences and workshops proceedings, book chapters, and technical reports, most notably [16, 17, 25, 26, 28, 29, 30, 33, 34, 35, 36], are presented in the three subsequent sections.

1.2.1 Scenario-Driven Prototyping Method

To avoid the development of irrelevant prototypes by engineers or a technologically unsophisticated studying of implications by social scientists, the two fractions obviously have to be closely coupled within an interdisciplinary process. Scenario-driven prototyping thus proposes the development of both scenario and prototypes emphasizing the main aspects of the scenario.

The process consists of four main steps: Firstly, in an interdisciplinary dialogue, the engineers present the possibilities of the technology, likely developments in the near future, expected costs, and so on. Based on this presentation and on the societal issues identified by social scientists (e.g., developments with societal conflict potential), they together search for relevant application domains and particular applications inside these domains. The result of this first step are future scenarios. In the second step, one or more prototypes are developed for each scenario, which underline the scenario's main aspects. Thirdly, new technological and societal issues are identified by analyzing scenarios and prototypes. The so-discovered issues can in a fourth step influence back the scenarios and subsequently the prototypes. Consequences with large societal conflict potential might for example be eliminated by technological means, while discovered technological problems might render parts of the scenarios unrealistic, which have to be rewritten.

More than unifying the specific advantages of prototyping and scenario-analysis, this "reality-check cycle" is the main synergistical advantage of the paradigm. While scenario-driven prototyping emphasizes from the outset the development of relevant, realistic scenarios and prototypes through the interdisciplinary process, it also forces through this feedback loop a mutual insemination between technological development and societal analysis – an insemination leading to a high degree of focus, realism, and relevance.

1.2.2 The Smart Tachograph

The Smart Tachograph is a first proof-of-concept for the scenario-driven prototyping paradigm. Within the project, we have first identified in a dialogue with economists that many pricing models might change from the nowadays static towards a dynamic and individual pricing, as soon as technology would enable it. Being recognized as a relevant domain, both economic and technologically realistic, we have further pursued it, writing a scenario.

Our original scenario envisioned dynamic and individual pricing in a supermarket, where products enabled with ubiquitous computing technology would change their prices according to supply and demand, identity and behavior of customers, and environmental context. Although there are good economic arguments in favor of such a model, and it would also be technologically possible some years from now, in a further round of interdisciplinary discussions, the scenario has been pointed out as unrealistic, due to the expected missing customer acceptance of fast changing supermarket prices.

After setting up new criteria for the likely acceptance of dynamic and individual pricing, we have found another, greatly more realistic domain for such model – vehicle insurances. The second, rewritten, scenario thus revolves around a system, the Smart Tachograph, that continuously measures the conditions and manner in which a vehicle is driven, and charges motorists on a pay-per-use/pay-per-risk basis the various types of costs they inflict. To make the prototype as close to reality as possible, we have not only developed the software for the system, but have also deployed real sensors to measure various driving parameters and continuously send the data to the computer analyzing them and transforming them into caused costs. Moreover, we have also developed the whole surrounding infrastructure which would be scalable to an arbitrary number of vehicles.

Following the scenario-driven prototyping method kept us from developing a prototype (and then analyze its implications) in an area which seemed at first realistic from several points of view, but has proven in the end not to be. It has, moreover, enabled several other relevant insights: We were able to discover some pragmatic technological limitations, which, although not of fundamental nature, render the measurement of some of the parameters influencing traffic costs unrealistic for the years to come. We have further discovered that although technologically and economically meaningful, and although generally accepted by the public, if measuring too many parameters, thus making the system too precise, a more subtle lack of acceptance would arise. Finally, and above all, by systematically taking into account technological means, economic drivers, and societal values, we have been able to prove that a fine-granular measurement and accounting of traffic costs is possible without the massive privacy intrusion characterizing similar proposals.

1.2.3 The Chatty Environment

The other application developed according to the scenario-driven prototyping method is the Chatty Environment, a prototypical mobile system designed to help the blind and visually impaired to lead a more independent life. When starting our project, several already existing research projects had the aim of providing the blind with navigational aids, usually on a GPS basis, similarly to car navigation systems. Setting out and discussing with blind persons their everyday problems, we have discovered, however, that there is another class of problems they regularly encounter – not being able to perceive their immediate surroundings. In a supermarket, for example, since all packed food feels the same, blind persons are at great trouble while shopping.

After realizing this, we have shifted the focus of the scenario and of the later prototype from helping the blind to physically navigate between two points, to a system that reveals them their immediate surroundings as they pass by, in a similar way as sighted people perceive their physical neighborhood. The system further enables them to logically navigate through the descriptions of the objects they pass, and even to electronically trigger actions in the real world, such as remotely opening a train's door.

To do so, we have equipped real-world entities with different types of electronic tags, which can be read by the blind users' mobile device when arriving in the neighborhood. We have also designed and developed a novel audio-input and press-select interface to the user. All objects picked up by the device are queued according to a priority system, and their textual descriptions are read one after the other to the user by a text-to-speech engine on the mobile device.

Similar to the Smart Tachograph, by following the scenario-driven prototyping process that now also from the beginning involved stakeholders, we could identify here a multitude of relevant issues as well, both on a technological and on a societal level: As already mentioned, we learned from the beginning that revealing the immediate surroundings to the visually impaired user is at least as important as navigating him or her to specific points, being thus able to propose for the first time such an application. After having the prototype ready, we conducted a second round of interviews with potential users, letting them experience the prototype. In this step, we were surprised to learn that there can almost never be too much audio information for the blind user. The users were usually happy with as much information as possible, telling us that information about the surroundings is what they often miss most. Thus, the complex information filtering algorithms we had originally planned became superfluous. We could further contribute with different insights about the design of an interface for the blind for the blind and visually impaired, which are likely to have an influence beyond our project. Finally, we have learned that the real-world context of such an application is so dynamic, that any system aiming at revealing the world to the user – and not necessarily for the blind, this could apply as well for tourists, people not speaking the local language, or children – cannot gain its information from a statical source. Relating a location to the context statically stored on a CD might work well for vehicle navigation systems, where the street topology is largely constant over time, but certainly not for such a system as the Chatty Environment. We have as a result proposed a novel class of context-aware systems.

1.3 Thesis Outline

In the second chapter, we provide an overview of the technological background in which this thesis has emerged. It starts with a historical perspective, presenting in more detail the vision of ubiquitous computing. The chapter then presents – from today's perspective – the technological drivers for ubiquitous computing, making Weiser's vision to seem more realistic than ever. To show how vast the application domain of ubiquitous computing can be, the chapter concludes by presenting a collection of typical applications, structured in application areas.

However, technological feasibility and various application areas alone do not yet induce the societal importance that we expect ubiquitous computing to have. For such applications to become reality, economic drivers must exist behind ubiquitous computing as well. We present them at the start of chapter 3. We further argue in the third chapter that, given the technological and economic drivers, together with the various application areas, we expect ubiquitous computing to have a strong societal impact. We thus dedicate the remainder of the chapter to present examples of societal opportunities and risks encompassed within the large-scale deployment of ubiquitous computing, motivating thus the importance of taking a closer look at them in a structured process of technology assessment.

After having motivated its importance, we start the fourth chapter by presenting the two approaches predominantly used so far in the assessment of societal impact of ubiquitous computing: technology-driven prototyping and scenario analysis. Scenario analysis being used in many domains as a systematic approach for understanding possible futures, we present its history in greater detail, before narrowing the discussion to the related work, consisting of other ubiquitous computing scenario analyzes. After closing the chapter by presenting the specific drawbacks of the two methods, we introduce in the fifth chapter our novel scenario-driven prototyping paradigm, and argue why besides unifying the individual advantages of the two classical approaches, it also brings specific features that could only have appeared from the synergy of the two. The chapter then provides a survey of the interdisciplinary technology assessment project in which the ideas presented in this thesis have appeared for the first time, and ends with an analysis of the limitations of our proposed method.

Chapters 6 and 7 are devoted to the in-depth presentation of the two projects in which we have applied the scenario-driven paradigm to different areas: personalized and behavior-dependent accounting of traffic costs, and the assistance of blind and visually impaired, respectively. The two chapters have similar structures. Following the method developed in chapter 5, they start by motivating why the particular domain chosen for scenario and prototype is societally relevant and well-suited for the analysis of ubiquitous computing consequences. Then, the chapters present the original scenario, its first analysis, and, if applicable, the rewritten scenario. The prototypes which we have developed to underline the main aspects of the scenarios are further presented in detail. Before concluding with the related work, both chapters present the actual aim behind the development of the corresponding scenario and prototype – the analysis of technological and societal consequences which they enable. These parts of the two chapters are the true measure for whether the method of scenario-driven prototyping has been a success or not.

2 Technology Background and Trends

This chapter presents the broader ubiquitous computing context in which this thesis emerged. It starts with a historical perspective, presenting the vision of ubiquitous computing, as envisioned by Mark Weiser and his fellow researchers at the Xerox PARC research laboratories. Then, the chapter presents several technological drivers for ubiquitous computing, making the case that taken together, the technologies comprised within the ubiquitous computing paradigm form a highly relevant technological topic for the near future. To give a flavor of the vast areas that could be affected by the deployment of ubiquitous computing technologies, several application domains of ubiquitous computing are introduced, along with highlighted individual applications. Given that over the last years several other terms, such as Pervasive Computing or Ambient Intelligence, have been introduced to denote a similar paradigm as ubiquitous computing, the chapter further introduces such terms, presenting the (large) similarities and (small) differences among them.

2.1 The Vision of Ubiquitous Computing

The term "ubiquitous computing," UbiComp for short, was coined during the late 1980s at the Xerox Palo Alto research center (PARC), by a group of researchers around Mark Weiser. They envisioned a future computing environment in which myriads of computers would dilute into the environment, people not noticing their existence any more, but only taking advantage of the services they bring to us in an unobtrusive way. The researchers started from the observation that many revolutionary technologies that changed people's lives in a sustainable way work in the background without being noticed: "The most profound technologies are those that disappear. They weave themselves into the fabric of everyday life until they are indistinguishable from it" [163]. Mark Weiser illustrates this statement with two such tools widely used by mankind: writing, which enables much of our modern communication, yet we usually only perceive the message, not the tool itself, and electrical engines. While during the industrial revolution, a factory typically had a single engine driving dozens of different machines through a complicated system of wires (see the model in Fig. 2.1), some hundred years later there are dozens of motors in such a cheap product (compared to an entire workshop building) as a car. These engines lift the windows and headlights, move the seats, lock and unlock the doors and so on, yet we usually do not perceive the small motors any more, but only the service they provide.

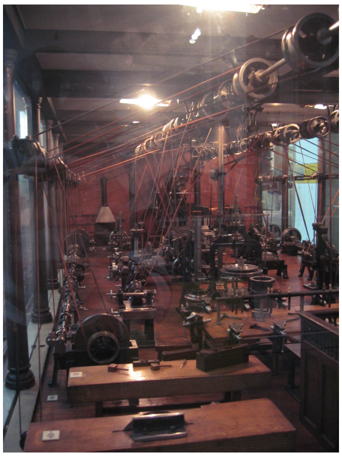

Figure 2.1: Model of an old workshop building as displayed in the Science Museum London.

In similar manner, Weiser and his team envisioned that computers are soon to disappear in the background, becoming "tools through which we work, and so which disappear from our awareness" [164]. They would be embedded in everyday objects such as mirrors, toothbrushes, pens, or even parking lots, and extend these artifacts with new, helpful functionality (e.g., the "smart" parking lot would know if it's free, publish this information, so that it can be picked up by some "smart" car, that would be able to reserve it and help the driver navigate to the free lot). We would interact with such embedded computing devices in the natural way we are used to interact with the objects they are embedded in. The computing-enabled artifacts would not only communicate with humans, but also among each other in order to accomplish their tasks (i.e., the vehicle communicating with the parking lot). Most importantly, they would do so unobtrusively and with no annoyances for the user – Weiser thus foresaw an age of *calm technology* [165].

Weiser's vision is not restricted to some specific classes of artifacts; his vision also goes well beyond electronic or electrical equipment only: "next generation computing environment in which each person is continually interacting with hundreds of nearby wirelessly interconnected computers" [164]. Virtually any real-world object could thus take part in the ubiquitous computing environment.[1] This holistic prediction is even more impressive if one thinks of the time it has been conceived: the late 1980's were the time of the Intel 386 generation of processors and Windows 3.11, it was the time when 14 inches CRT-displays were in use, bulky mobile phones were a costly gadget for a few highly-ranked decision makers, and the World Wide Web still had a few years before starting to cover the world with information.

2.2 Technology Drivers for Ubiquitous Computing

While Weiser's ideas may have seemed pretty visionary more than ten years ago, the rapid technological advances over the last decade make his visions seem much more feasible today. Mattern, professor at ETH Zurich's Institute for Pervasive Computing, for example, identifies four major technological trends, which, taken together, will allow more and more everyday objects and physical environments to become "smart," guiding our society towards Weiser's vision [111].

The *first* trend is the ongoing miniaturization in the field of microelectronics. Gordon Moore predicted in the mid-1960's what has become ever since known as Moore's Law – that the computational power of microprocessors per surface unit will keep doubling every 18 months, as had already happened for ten years before [117]. Not only has this "law" held true with an astonishing precision for roughly 40 years since – later on, it has also been successfully extended to storage capacity and communication bandwidth. At the same rate as the miniaturization of electronic components progressed, the prices for those components kept falling. After half a century of exponential decrease in both size and price, we have arrived at a point where processing power and memory are cheap and small enough to be built into almost any everyday item without substantially increasing its price or dimensions. And this exponential trend is expected to last for at least another 15 years [110].

---

[1] A few years later, Weiser wrote: "The next generation Internet protocol, IPv6, can address more than a thousand devices for every atom on the earth's surface. We will need them all." [165]

The *second* driving factor is formed by the rapid progress in wireless communication technologies. Not only is the bandwidth steadily increasing; several new standards such as WiFi, Bluetooth,[2] ZigBee,[3] or NFC (near field communication)[4] cover a broad application domain, allowing smart objects to interconnect over shorter or longer distances, depending on the application's needs.

Together with the *third* trend towards better and smaller sensors (some even functioning autonomously without batteries, e.g., by harvesting the needed energy from environmental temperature changes), all these technological advances show a clear tendency towards the vision of smart objects and smart environments, that are able to feel, think, memorize, and talk. Additionally to these broad technological trends (of which the fourth will be presented below), Mattern also identifies two particular technologies that can be understood as special and outstandingly important kind of sensors [112]. First, identification technologies such as RFID (radio frequency identification) allow for an automatic identification of entities (i.e., objects or people). Automatic identification of more and more real-world entities means a further leap towards an infrastructure of smart things and smart environments. Second, location technologies can also be seen as a special class of sensors, which provide location information. Global positioning systems such as GPS[5] or the emerging European Galileo system[6] are only the tip of the iceberg. Not yet as widely perceived by the public, indoor-location technologies have also reached an unprecedented level of accuracy. The still relatively expensive ultra-wideband (UWB) technology[7] is able to provide indoor location information with a precision of a few dozen centimeters. Other researchers have suggested the use of signal strength from wireless communication technologies such as WiFi and Bluetooth to gain indoor location information – not as precise as with dedicated systems but at virtually no additional costs [18]. And there are numerous examples of rather exotic location technologies, such as using the signal strength of FM radio stations in an urban area [93], or using smart sensing carpets for indoor location [58].

Taken together, these enabling technologies show a clear trend to-

---

[2] See www.bluetooth.com.

[3] See www.zigbee.org.

[4] See www.nfc-forum.org.

[5] See http://waas.stanford.edu/.

[6] See http://ec.europa.eu/dgs/energy_transport/galileo/index_en.htm.

[7] See, for example, www.ubisense.net.

wards the “Internet of things,” as envisioned by Gershenfeld [62] and thoroughly regarded by Mattern [112]. However, the Internet of things is usually envisioned not as an end it itself, but merely as means to realize the holistic vision of ubiquitous computing. Humans will interact and take advantage of the services provided by the “Internet of things.” Since the myriads of computers then surrounding us are to bring their services and help us in an unobtrusive way, novel interfaces that go far beyond the classical input and output devices such as keyboard, mouse, display, and printer have to emerge.

Such interfaces will benefit from the *fourth* technological trend stressed by Mattern [111], the recent progresses in material sciences. They led to the development of entirely new materials, which will allow computers to come in wholly new forms and people to interact with them in entirely new ways. Light-emitting polymers, for example, can be used to create flexible displays, which can then be integrated into arbitrary objects such as windshields, milk cartons, or cereal boxes [99]. Prototypical displays using “organic LEDs (light emitting diode),” have been presented over the last years and will most likely soon be available for commercial use [166]. “Smart paper” using “smart ink” is another example of an upcoming ubiquitous computing technology that could offer a novel interface to the digital world. Smart ink consists of microcapsules in which white and black pigments, differently charged, flow freely. By applying a positive or negative voltage to a point of the smart paper covered with smart ink, that point changes its color [44]. Since a voltage must be applied only to change the color, not to maintain the new one, the voltage can be applied either by an electrically charged pen, or by the underlying surface. Thus, smart paper can be used as both input and output medium, the advantages of a computer display being combined with the natural way of using a piece of paper [112]. As can be seen from this example, such novel ubiquitous computing interfaces have the potential to free the user from concentrating on the interaction with computerized systems. The user is able to focus on the task by using a well-known interface, while still taking advantage of the underlying computerized systems.

Eyeglasses that can project – by a small laser that is almost invisible to others – images directly onto the retina, are another new way of interacting with computers. They have been already around for some years now and could enable completely new applications. Half-jokingly, Satyanarayanan – professor at Carnegie Mellon University and founding editor-in-chief of the IEEE Pervasive Computing journal – described in an interview such possible usage scenarios: “You could wear a pair of glasses with a small amount of face recognition built-in, look at a person, and his name would pop up in a balloon above his head. You could know instantly who the person is, even if you don’t immediately recognize him. I look at my tree, and a little balloon pops up saying, ‘Water me,’ I look at my dog, it says, ‘Take me out,’ or I look at my wife, it says, ‘Don’t forget my birthday!’ ”[137]. Such novel ways of interaction with computers could not only open numerous new windows to the virtual world, they could also have a large effect on how people interact with each other. Projects such as the “Lover’s Cups” [23] show how peripheral and pretty subtle channels that go far beyond the classical e-mail or instant messaging paradigms could help friends living at a distance in having a sense of what the other is doing – in this special case, when the far-away partner is drinking or how full his or her cup is.

Taking a more user-centric approach, Davies and Gellersen identified two major drivers towards the public acceptance of ubiquitous computing. They argue that in the years that have passed since Mark Weiser came up with the vision of ubiquitous computing, the society at large has undergone two major paradigm shifts, which should contribute to a much easier acceptance of pervasive computing technology than a decade ago [39]. One of these paradigm shifts has been the rapidly-paced wide-spreading of mobile phones in the 1990’s. Due to the fact that almost anyone owns a mobile phone today combined with the major drop in prices for long-distance telephony, people became used to be unobtrusively helped by technology to stay connected with loved ones and business partners almost regardless of their geographical location. The other paradigm shift has been introduced by the emergence of the World Wide Web over the same last decade of the 20th century. Through the WWW, so Davies and Gellersen, people further became used to be able to instantaneously access countless information sources. Due to matured search engines they can do so with very low transaction costs (i.e., costs to find information). Taken together, these paradigm shifts from the 1990’s should ease today’s acceptance of ubiquitous computing applications, which could be perceived by the public as further steps along the trend of being ubiquitously connected to information and people.

2.3 Ubiquitous Computing Applications

After having seen which technology drivers point towards a rapid widespreading of ubiquitous computing applications, and why these are likely to be accepted by the public, this section presents an overview of early, current, and likely future applications of the ubiquitous computing paradigm. This will give a flavor of how broad the possible application domains for ubiquitous computing are; motivating thus the importance of looking at its possible societal consequences – a task that will be accomplished over the next chapters.

2.3.1 Early Prototypes

In the early years of ubiquitous computing, researchers at the Xerox research center and in other laboratories started to built first prototypes. At Xerox PARC, Weiser and his fellow researchers experimented especially with new forms of interaction between computing systems and people, devices that go far beyond the classical desktop computing paradigm. They developed "inch-scale machines that approximate active Post-It notes, foot-scale ones that behave something like a sheet of paper (or a book or a magazine), and yard-scale displays that are the equivalent of a blackboard or bulletin board" [163]. In the early 1990's, at the Olivetti research labs in Cambridge, UK, a team around Roy Want and Andy Hopper kicked off another area that later on proved to be of utmost importance for the field: *location awareness*, a term that refers to (mainly mobile) applications that behave differently depending on the location of users and/or other entities. They developed and deployed the so-called "active badges," electronic tags used to locate their wearers through a sensor network infrastructure receiving the pulses periodically sent by the tags [162]. The active badges proved to be a classical enabling technology, whose whole potential only became clear after building the prototype. Several previously unanticipated uses of the system emerged when all lab members started wearing the active badges and could thus be located, such as the secretary being able to route incoming calls to the correct rooms (or not to route the calls if the presence of several lab members indicated a meeting taking place in that room). One other prominent research facility is the "things that think" (TTT) consortium founded in 1995 by the MIT media labs.[8]

---

[8] See http://ttt.media.mit.edu/.

Neil Gershenfeld, at the time director of the TTT-consortium, once wrote his credo, which resembles quite a lot Weiser's vision: "In retrospect it looks like the rapid growth of the World Wide Web may have been just the trigger charge that is now setting off the real explosion, as things start to use the Net so that people don't need to" [62].

Within the German-speaking countries, the Telecooperation Office (TecO for short) in Karlsruhe has had an important impact on ubiquitous computing research. In the late 1990's, the group around Hans-Werner Gellersen, Albrecht Schmidt, and Michael Beigl built the MediaCup, one of the first examples of everyday objects made "smart" through the use of ubiquitous computing technology [60]. The MediaCup is a coffee mug enhanced with several sensors, such as acceleration, pressure and temperature. The cup thus gains some basic kind of "self-consciousness", knowing the amount and temperature of its content (i.e., coffee). It also has wireless communication capabilities, being thus able to connect to other cups, the coffee machine, or other smart objects. The TecO and Hans-Werner Gellersen have also been the first to organize a conference dedicated exclusively to the field. The "International Symposium on Handheld and Ubiquitous Computing" (HUC'99)[9] organized in Karlsruhe in 1999 has not only been the first conference focusing exclusively on ubiquitous computing, it is still the largest and most important academic conference in the field, carrying the new name of "UbiComp".[10]

2.3.2 Countless Possible Applications

From such early prototypes, however, the field of ubiquitous computing has grown explosively over the last decade or so, encompassing numerous technologies, and covering nowadays virtually every field of human activity as possible application field.

Logistics

From the very beginning, one of the driving forces behind ubiquitous computing have been logistics. Ubiquitous computing technology, such as RFID tags, enables companies to track their assets in real-time, and thus to improve various economic parameters, such as inventory management, supply chain efficiency, product availability, while at the

---

[9] See www.teco.edu/huc/.

[10] See http://ubicomp.org.

same time avoiding typical negative phenomenons such as the bullwhip effect [52]. Moreover, by tracking not only the location, but also the status of their products, companies can further increase their efficiency. Airplane turbine manufacturers, for example, equip their products with sensors that continuously monitor the status of both critical and non-critical systems. If such a "smart" engine detects a failure of a non-critical system during flight, it automatically orders the spare part at the destination airport, which can thus be replaced right after landing, minimizing the non-productive grounding time of the plane [141].

Public and Private Transport

Several other instances of both private and public transportation are important application fields of ubiquitous computing. One reason has been pointed out in [29]: vehicles are large and expensive enough to allow some ubiquitous computing technology to be built in without a major impact on neither size nor price. Furthermore, they also provide enough energy to power such systems. These arguments certainly apply to public transportation – such as trains, or the bus fleet of a large city – as well.

Reaching further than mere navigation systems (which themselves can be considered a simple form of ubiquitous computing systems), "smarter" cars have started to be equipped in Japan as early as 1997 with a more intelligent navigation system, which also takes into account the momentary traffic situation to determine the optimal route [66]. With gradually improving technology, many new services will become feasible. A good overview is given by [70]. For example, having both a robust method of communication between trucks and a dependable way of measuring the distance to the truck in front could let trucks on a highway form an assembly with very small distances between them, improving the fuel efficiency of every single truck in the assembly. If one day a positioning system with an accuracy of a few centimeters would become available, every movement of a car could be determined in real-time, related to the movements of surrounding cars, and in case of an imminent danger, automatic systems could gain control over the vehicles in order to actively avoid collisions. Public transportation could take advantage of ubiquitous computing technologies as well. One example is pointed out in [30]: if smart trains know the final destination of all their passengers, in case of delays a much better decision can be taken on whether other trains should wait or not, depending on the number of passengers that have to transfer and their alternatives. Nowadays, such decisions are taken by the railway staff based on their experience, which often leads to suboptimal solutions and increased dissatisfaction among travelers.

Public Services

Several other public services could be application fields for ubiquitous computing. In his founding article [163], Mark Weiser uses such a public service to underline the possible uses of ubiquitous computing: While Sal, the protagonist of Weiser's scenario, is approaching the office, her smart car starts searching for smart parking lots that have published to be free, and subsequently reserves one of them and guides Sal to it. Electronic guides for tourists [37] or for visually impaired [33], which use location information to provide information about the neighborhoods to their users, might be other examples.

Working Environment

Future offices could also display a wide range of such technologies. One of the main uses of the early Olivetti active badges [162] was for the telephonist to know in which room the employees are situated and to thus be able to easily route calls, but also to conclude that there's a meeting going on if several people are in the same room and not forward any calls in such a case. In later projects, smart rooms would autonomously derive whether there is a meeting going on (for example, by noticing a gathering of smart coffee cups [61], or from the events generated by smart chairs that are equipped with weight sensors [11]). After inferring an ongoing meeting, such smart rooms would then modify their own behavior [30], or that of other entities (for example, turning the mobile phones of participants to silent mode [61]).

Household Appliances

Many household appliances could take advantage of ubiquitous computing technology as well. A smart fridge[11] could be programmed with the allergies of the inhabitants. It would check the ingredient list every time a new product is deposited inside (by communicating with the smart

---

[11] Which should not to be mistaken for the over-popularized "intelligent," Internet-connected fridge, whose "intelligent" behavior is limited to automatically reordering exhausted products, a functionality that would quickly become annoying [27].

products) and issue a warning in case of incompatibilities [27]. Similarly, a smart washing machine would not need to be programmed, as it could directly ask the laundry for washing instructions [17]. A smart medicine cabinet, already existing as prototype, would be connected to the family doctor's office, automatically receiving the prescriptions for all family members. It would issues warnings when a non-prescribed medicine is taken, reminders when someone forgets to take prescribed medicine, and would even know in real-time about a medicine that has been withdrawn from the market because novel data shows it could be dangerous, and subsequently warn the patients of the danger [54].

Healthcare and Assistance for the Disabled

Going one step further than the medicine cabinet (which can only know whether a pack of medicine has been taken out or not, not if the medicine has subsequently also been ingested by the patient), a smart blister pack knows when a single pill has been taken out and reports this to the medicine cabinet [56].

In HP's groundbreaking "CoolTown" project[12] one of the visions referred to patients with a health condition that does not require continuous treatment but continuous observation, due for example to the possibility of a sudden and serious degradation of their condition, case in which they would need immediate treatment. Nowadays, such patients are typically held in hospitals. By monitoring their parameters and notifying the hospital in case of an emergency, such patients could lead a perfectly normal life outside hospitals and still receive the needed emergency treatment on time, when necessary.

Another group of patients that could gain more independence from hospitals and nurses through the use of ubiquitous computing are elderly patients suffering of milder forms of dementia. Several projects address such issues, e.g., by using cameras and activity recognition software to automatically notice if the users have performed such actions as brushing teeth and reminding them in case they have not [6, 116].

In a different healthcare domain, people with physical or mental handicaps could also take advantage by either gaining more everyday in-

---

[12] The project being already finalized for a few years, its former webpage, www.cooltown.hp.com, seems to have been permanently removed. The page contained professionally done videos depicting the visions behind the project, visions that have not been published otherwise. Due to the large impact of the CoolTown project, however, these visions have become basic common knowledge within the ubiquitous computing community. [87] describes more the technical aspects of the project.

dependence or be better included into their social environment. Several projects have aimed at guidance systems for the blind, to enable them to independently navigate through unknown places (e.g., [68, 123]). Other projects use a different approach, revealing the immediate environment to the visually impaired, letting them interact with them, and thus gain a quick understanding of their surroundings [25].

[15] proposes the use of ubiquitous computing technology to enhance a jigsaw puzzle with smartness, so that it can give hints to the players, thus allowing people with cognitive disabilities to take part in the game. [55] proposes the use of similar technology for playing card games, where individual players might receive the aid of personal electronic assistants. This would allow novice players (e.g., children), cognitively disabled persons, but also visually impaired to be more quickly and better integrated in such a game, without diminishing the game pleasure for the other players.

2.3.3 Sensor Networks and Smart Environments

A related direction of research are the so-called "wireless sensor networks" [131]. A wireless sensor network consists of numerous sensor nodes, which are autonomous small computation and communication devices equipped with sensors. Typically, such sensor nodes would be spread in the environment, where they would measure environmental parameters, and collaboratively perform a task, such as transporting the gained information to some central computer, or trigger an alarm upon cooperatively deciding that a certain temperature threshold has been trespassed. This cooperation among sensors constitutes the sensor network. Popular visions include using sensor networks for monitoring animals in the wild, or using them to monitor certain parameters in buildings, such as the temperature. The "smart dust" project [86] had the goal to take the paradigm of sensor networks to the extreme and build extremely small, dust-size (i.e., sub-millimeter) sensor nodes. Although the project never reached its intended goal, it came up with some interesting solutions and showed where the journey may lead in several years. Sensor networks highlight a slightly different part of ubiquitous computing than the above-mentioned smart objects. Sensor networks are typically not concerned with enhancing objects, but with digitally augmenting the physical environment itself – the term of "smart environment" is often used to describe the result.

2.4 Terminology

After this survey, two questions are still open: what exactly does "ubiquitous computing" mean, and what is the difference to related terms such as "pervasive computing" or "ambient intelligence"? As a start, it is important to notice that the term "ubiquitous computing" is used with several meanings. Ubiquitous computing surely stays for the *vision* put forward by Mark Weiser, of a "next generation computing environment in which each person is continually interacting with hundreds of nearby wirelessly interconnected computers" [164]. But it also describes the *collection of technologies* that enable this very vision, some of which have been presented earlier in this chapter. Further, ubiquitous computing is often used to denote the new paradigm of such novel computing environments as a *goal* that is worth pursuing. Finally, ubiquitous computing denotes an *area of research*. However, the ubiquitous computing research does not only comprise the enabling technologies themselves. It is also concerned with the applications that may arise by using those technologies, with the novel user interfaces needed to interact with disappearing computers, and with the strong societal, privacy-related, and economic consequences of a paradigm that intends to influence numerous aspects of everyday life and needs to heavily monitor users in order to do so.

Other, more recent terms are often used to describe technologies and visions similar to those of ubiquitous computing. "Pervasive computing" is a more industry-close term taking a bottom-up approach that starts at technologies available today or in the near future (compared to the more holistic, top-down vision of ubiquitous computing), while at the same time placing a greater weight on the business cases related to those technologies. The pervasive computing vision can probably be best summarized with Lou Gerstner's words from 1997, chairman of IBM at that time: "A billion people interacting with a million e-businesses through a trillion interconnected intelligent devices". The European Union's "Information Society Technologies Program Advisory Group" (ISTAG) introduced in its 1999 vision statement the term "ambient intelligence" to describe a vision where "people will be surrounded by intelligent and intuitive interfaces embedded in everyday objects around us and an environment recognizing and responding to the presence of individuals in an invisible way" [3]. "Ambient intelligence" thus emphasizes more explicitly the human-centric part of the vision and the therefore needed novel interfaces. For the purpose of this thesis, the nuances of “ubiquitous computing,” “pervasive computing,” and “ambient intelligence” will be let aside and the terms used interchangeably. Sensor networks are considered a specific sub-domain of ubiquitous computing research, while “smart environments” can be both the concrete result of ubiquitous computing and sensor networks projects, and a similar vision to ambient intelligence, and thus to ubiquitous computing.

2.5 Summary

The aim of this thesis is to provide a novel approach to realistically assess the consequences of the envisioned wide-scale deployment of ubiquitous computing technologies in the near future. This technological background chapter laid the foundation for the argument that we will develop in the following chapters. After introducing the vision of ubiquitous computing, the chapter argued why this vision seems realistic. It presented a collection of emerging technologies which, taken together, enable the vision to become reality. A wide range of applications building on these technologies have then been presented, illustrating the large possibilities encompassed within pervasive computing.

Thus, the current chapter had two aims: First, it tried to show that the vision of ubiquitous computing is *likely* to become reality, at least from a technological perspective. Second, it made the case that *various aspects* of everyday life could take advantage of the technological progress.

The next chapter will build on this argument. It will show that apart from the technological drivers behind ubiquitous computing, there are also strong economic interests and social drivers sustaining its deployment. Technological, societal, and economic drivers taken together make a strong case for the expectation that ubiquitous computing technologies will be deployed on a large scale in the near future, bringing numerous benefits, some of which have already been presented. However, as with probably all technology trends likely to have an important impact in the future, there is not only a bright side of the story. Ubiquitous computing also holds several dangers, which will be presented within the next chapter as well.

The technological likelihood, together with the economic and societal drivers portrayed in the next chapter, make the use of ubiquitous computing technologies in the future to seem quite probable. Under this premise, and given that the range of feasible societal gains, but also of possible dangers resulting from ubiquitous computing is quite large, the necessity of systematically analyzing these societal consequence at an early stage of technology development becomes evident. How such analyzes have been done in the past and what our approach is, will be the topic from chapter 4 onwards.

3 Societal Background and Implications of Ubiquitous Computing

By looking at the powerful technology drivers in section 2.2 and the many application areas from section 2.3, it becomes evident that ubiquitous computing is likely to have a profound impact on many areas of our daily lives, and on our societies in general. This chapter gives a first overview of the societal opportunities and risks that are to be expected through the spreading of ubiquitous computing technology.

It starts by presenting the deep impact that ubiquitous computing could have on several business areas. This represents both a non-technological driver for ubiquitous computing technology (through the promise of increasing their efficiency or the perspective of completely new businesses, companies are likely to adopt the technologies enabling them for these advances), but also a likely societal consequence (the modified or entirely new business processes changing to a certain degree the society in turn – in both positive and negative ways).

Subsequently, the chapter presents a great deal of other societal risks and opportunities of ubiquitous computing technologies. Especially potential risks of ubiquitous and pervasive computing have been broadly and often rather critically analyzed by scientists from different fields – section 3.3 gives an overview.

The chapter concludes by making the case for the difficulty of a reasonably realistic assessment of the ubiquitous computing implications. Obviously, every future technology assessment is difficult. However, for ubiquitous computing, it means considering a vast collection of technologies that could influence virtually any aspect of everyday life, bringing along a remarkably heterogenous collection of both opportunities and pitfalls. This will serve as motivation for our approach of scenario-driven prototyping for ubiquitous computing, that will be presented in chapter 5.

3.1 Economic Drivers for Ubiquitous Computing[1]

Section 2.3 has already presented some of the existing and envisioned application areas for ubiquitous computing technology. Through the vast spectrum of such areas, the expected importance of ubiquitous computing has been stressed. However, section 2.3 has focused only on technological feasibility and has not been concerned with the economic and societal premises needed for such scenarios to become reality.

Obviously, the mere technological feasibility does not imply that a scenario will be realized. At least one other condition has to be fulfilled:[2] An organization must be willing to invest in the development of the respective application. This could either be the industry (which would only do so when expecting a good business model behind), or some public authority (which could do so for several other reasons). Thus, in contrast to section 2.3, this section explores the benefits that are to be expected through the deployment of ubiquitous computing from a business perspective. It shows how companies could increase the efficiency of existing business processes or come up with entirely new business models through the use of ubiquitous computing.

Hence, this section serves as motivation for the importance of analyzing the social impact of ubiquitous computing. Since – as argued above – mere technological feasibility does not imply the realization of a scenario by itself, the economic ubiquitous computing drivers presented here round off the picture started with the technological drivers in chapter 2.

3.1.1 Extensive, Ubiquitous Information

The industry has already begun setting its sights on the enormous business potential that technologies such as wireless sensors, RFID tags, and positioning systems have to offer. Analysts call it the "real-time economy" or "now-economy" [141], where more and more entities in the economic process, such as goods, factories, and vehicles, are being enhanced with comprehensive methods of monitoring and information

---

[1] This section is based on joint work with Jürgen Bohn, Marc Langheinrich, Friedemann Mattern, and Michael Rohs [16, 17].

[2] Another condition that must be met is the societal acceptance of the respective application. This, however, is a more fuzzy domain, and often hardly to grasp, as will be argued throughout this thesis. It is precisely this difficulty of assessing the future impact of ubiquitous computing technologies, combined with the importance of doing so, which is one of the main motivations behind this thesis.

extraction. Ultimately, the whole lifecycle of products, beginning with the "birth" of their components and ending with their final disposal (or recycling), can be witnessed (and, to some extent, even controlled) in real time.

Two important technologies form the core of these new economic processes and applications: the ability to track real-world entities, and the introspection capabilities of smart objects. Tracking objects in real-time allows for more efficient business processes, while objects that can monitor their own status via embedded sensors allow for a range of innovative business models.

Increased transparency can be a worthwhile investment due to improved inventory management alone, which obviously benefits from accurate, real-time information on the location and condition of goods, equipment, and manpower. If a company does not know the location and condition of its stock, and how long it has been in the warehouse, significant costs are incurred. Missed profits, oversized inventories, and the devaluation of goods depreciating in the warehouse are possible consequences of a lack of information. The stocktaking required for business or legal reasons also typically requires a considerable amount of effort. Stocktaking is not only expensive; it is inherently error-prone as well. A factory floor or warehouse equipped with technologies such as indoor localization and automatic identification can largely automate the stocktaking process, thereby reducing costs.

If several companies along a supply chain simultaneously use such precise inventory data in addition to real-time order information, they can achieve additional savings by significantly attenuating the so-called "bullwhip effect" [100]. This effect, often noticed in practice, describes the following phenomenon: although consumer demand for a product remains almost constant over time, small changes in this demand amplify along the supply chain and ultimately result in either excess production (and associated storage costs) or sudden interruptions to supply (and associated missed sales). However, the more information transparency there is along the supply chain, the more these undesirable effects can be minimized [83]. By making comprehensive information available along the supply chain, a significant reduction in the bullwhip effect can be achieved.

A further step towards the now economy is the constant monitoring of critical product parameters (e.g., of temperature-sensitive goods such as chemicals or groceries) by tiny wireless sensors. Equipped with communications capabilities, such “introspective” goods are not only able to monitor themselves, but can also communicate relevant parameters to the outside world [51]. Smart goods could observe their condition while in transit and trigger an alarm in the event of excessive temperatures, which could – if appropriate – lead to an automatic reordering of damaged goods. Alternatively, the goods could also attempt to take corrective action, for example by controlling the temperature of their container: “As sensors improve and always-on connectivity becomes a reality, products will be able to do something about their condition” [49]. In this way, “self-conscious” products (i.e., products that perceive their condition, analyze it, and attempt to change their situation if they are dissatisfied with it) would lower the total transaction costs by reducing the time necessary to procure replacements for damaged goods.

3.1.2 Innovative Shopping Models

The benefits of a world full of smart objects do not stop at the factory floor. Once comprehensive infrastructures for tracking goods and facilitating communication among “self-conscious” objects are in place, ambient intelligence throughout our environment – in private homes, cars, trains, and public places – would facilitate a range of new applications and business models. The prospect of their revenue streams might become another substantial driver for the deployment of pervasive-computing technologies in the near future.

Real-Time Shopping

“See a great sweater on someone walking by? Find out the brand and price, and place an order. Or maybe you’ll be wearing the sweater and earning a commission every time someone near you sees and buys.” This vision [49] describes a future, maybe not that far off, in which the boundary between the real world and the world of information has become blurred, and people have become media. Invisible tags embedded in most products would allow consumer devices to read out the unique identification number of an item and use this to access the object’s virtual representation in the information world, which in turn could provide the user with a wealth of background information (e.g., ingredients, product reviews, etc.) as well as direct links to online (or even real-world) shops selling this item. People could shop on the move

– on the streets, in buses, or whilst chilling out in their favorite bar at night – and every item sold could in turn become a new sales channel.

In a more proactive fashion, smart products could begin to subtly advertise themselves, or even use cross-marketing to advertise their "friends." A smart refrigerator could, for example, recommend recipes based on both the groceries it contains, and on the items currently discounted at the local supermarket. It could also accrue reward points every time goods of a promoted brand are stored in it.

Highly Customized Prices

With the help of ubiquitous computing technology and the detailed profiles it enables, prices of everyday goods could even be adjusted to suit individual customers. In the vendor's ideal world "each consumer would be quoted an individual price, which [. . .] exactly corresponded to his readiness to pay" [143] – economists call this "perfect price discrimination."

More generally, a "perfectly competitive market" (for short: perfect market) is defined in classic economic theory [92] as having three characteristics: the traded goods are homogenous, the existence of complete and correct market information on both supply and demand side, and the equality of buyers and sellers in terms of space and time distance to the marketplace. Nowadays, only stock-markets can be regarded as perfect markets. Only here goods are perfectly homogenous (one stock being as good as the other), all players are concentrated in one place at the same time, and all have complete information about prices being asked and offered. Advantages of the perfect market include peak trading of goods and optimum price for a majority of sellers and buyers. Such markets are characterized by a highly-dynamic price structure. There is a permanent bargaining to find the current market price; based on this price market players come to a decision for their next-moment moves, which in turn determine next moment's price and so on.

As [98] argues, ubiquitous computing has the potential to transform many traditional, static marketplaces into such highly dynamic, perfect markets. One extreme example could be supermarkets. If smart products in a supermarket were able to sense their environment – e.g., other products around, the time of day, the day in the week, and so on – and communicate with other products, with the shelves and maybe the cash register, then supermarkets can become perfectly competitive markets where prices are determined dynamically, in correspondence to supply and demand. In short, “the supermarket becomes a stock market” [98].

Milk bottles, for example, could set their prices dynamically, according to parameters such as time of the day, day of the week, season, outside weather, the presence of other milk bottles in the neighborhood, and the own best-before date. When approaching this expiration date, a bottle would decrease its individual price, so customers will be tempted to buy it and not grasp for another, fresher one. The same would happen when sales decrease due for example to rain, and part of the stock risks to expire. On the other hand, when sales are exceeding expectations, bottles would increase their prices, as long as people are buying. Thus, a highly dynamic reaction to market facts is possible.

Aside of the questionable economical feasibility of such a model,[3] however, individual prices in general are to be regarded critically due to the reaction of customers. The online-bookseller Amazon, for example, discovered that such price “perfection” might not be universally appreciated, after their trial of individual DVD prices had to be suspended almost immediately due to massive customer criticism [141].

Objects Acting Economically Autonomous

The ultimate in shopping may be achieved when all the decision-making is removed from people, and things do the shopping themselves. The business consultancy Accenture has coined the phrase “silent commerce” for this. Their vision of “autonomous purchasing objects” not only includes photocopiers responsible for ordering their own paper, but also Barbie dolls that delight children (and their parents...) by ordering new clothes with their own pocket money: “Barbie detects the presence of clothing and compares it with her existing wardrobe – after all, how many tennis outfits does a doll need? The toy can buy straight from the manufacturer via the wireless connection... She can be constantly and anonymously shopping, even though the owner might not know it” [108].

3.1.3 Generalized Pay-per-use

In a future full of smart objects, we may not only be tempted to buy (and have the ability to shop) just about anywhere at any time, we

---

[3] A much more realistic example for individual pricing – individualized vehicle insurances – will be introduced in chapter 5 and have the entire chapter 6 devoted to.

may also *have* to buy just about everywhere, all the time. Digital rights management systems that have recently been developed for distributing digital music and video are a first step in this direction. These systems make it possible for owners of digital content to exert control over the access to digital information even after it has been sold. Digital rights management systems could even be programmed to require a continuous payment while listening to digital music or watching digital video. Such pay-per-use models have traditionally been used for phone calls or public utilities (e.g., electricity, gas), but could now – in an ambient intelligence world – be implemented using everyday objects equipped with sensors and communications capabilities. Furniture, for instance, could monitor its usage (e.g., a sofa could count the number of persons that sit on it, the persons' weight and seating time) and create a monthly itemized billing statement. While private homes might still prefer owning their beds rather then being billed for sleeping in them, corporate buyers with a high turnover of furniture (such as hotels or offices) could potentially provide a sufficiently large customer base for such a business model to develop.

3.2 Societal Opportunities of Ubiquitous Computing

After having looked at some economic drivers for ubiquitous computing, this section is concerned with the societal benefits that could emerge from the economically-driven and technologically feasible ubiquitous computing applications presented so far.

Behind merely improving the efficiency of businesses, the vision of a future filled with smart and interacting everyday objects offers a whole range of fascinating possibilities. Some of the main societal benefits that will presumably arise from the deployment of ubiquitous computing technologies are: better usage of energy and natural resources, increased security, improved healthcare and integration of the disabled, and new means of societal interaction across ages and over distances.

3.2.1 Energy Efficiency

In the ongoing international debate about climate change and ways of slowing down this process, a particular importance is given to means of increasing the energy efficiency of our modern societies. Smart environments and smart, “self-aware” objects promise to contribute to this process by providing the means to reduce the waste of energy on several levels.

Traffic is a good example for the energy saving potential of ubiquitous computing. Smart navigation systems that take into account the current traffic situation and route vehicles around traffic jams [66] would contribute to the reduction of greenhouse gases by minimizing the fuel wasted in traffic jams. The vehicles would not only be consumers of traffic jam information, they would also produce this information themselves – one car stopping on a highway could be broken, several vehicles stopping close to each other give a clear indication for a jam. Another typical waste of fuel occurs during the search for a free parking lot in crowded city centers. As already envisioned by Mark Weiser [163], car manufacturers have recently started to build prototypes of networked vehicles that inform each other over parking lots they have just cleared. Finally, a tight assembly of trucks on the highway (where the distances between the individual trucks measure only a few dozen centimeters), would increase the fuel efficiency of every single truck [70].

Garbage recycling might be another example for a near future in which ubiquitous computing could be used for a better exploitation of resources. By having more and more products tagged with passive RFID tags, the disposed products could automatically be sorted for recycling according to their constituents [13].

Moreover, besides these indirect positive effects through less waste of energy, ubiquitous computing could also contribute directly to the protection of the environment. Tiny communicating sensors the size of dust particles could be spread in the environment and detect, for example, the dispersion of oil spills or forest fires [17]. Wireless sensor networks have already being used for a large scale birds habitat monitoring application [153].

3.2.2 Security & Safety

One of the core features of ubiquitous computing is its potential for surveillance. Some higher-priced vehicles, for example, already come with the capability of noticing when they have been stolen (since they move without having read the RFID tag contained in the ignition key). In such a case, they notice either the owner or the police (also communicating their position acquired from a hidden GPS module). Moreover, they can even take action, by shutting the engine down.

But not only things can be tagged and tracked, similar technologies can be used for keeping track of people, too. Crimes or criminals can be detected more easily through a spreading mesh of technologies such as surveillance cameras throughout cities, automatic recognition of license plates (and thus, stolen cars), or biometric passports that are almost forgery-proof. Other, more subtle technologies, and with a narrower scope could be used, too. Parents would no longer lose track of their children, even in the busiest of crowds, when location sensors and communications modules are sewn into their clothes [17]. Or, from the comfort of their home, they will always know where the children are, having the location of their mobile phones tracked. Dementia patients could be tracked in similar manner, an automatic system in a smart home detecting any unusual activities and, if necessary, alarming paramedics [116].

Road safety has been continuously improved over the last decades through the use of computers embedded in vehicles. The anti-lock braking system (ABS) or the electronic stability control (ESC) are examples of two such features that have greatly improved the controllability of vehicles. In a slightly more distant future, smart cars that know their speed, direction, and their position with a precision of centimeters, will be able to avoid many accidents by exchanging this data and coordinating collision avoidance actions [70].

Another area where ubiquitous computing holds large promises is the early catastrophe recognition and the coordination and support of rescue forces in case of calamities. A tsunami early detection system is in a broader sense an example of ubiquitous computing technology [31]. The CodeBlue project from Harvard University aims at guiding rescue forces quicker to victims and coordinate their efforts for improved time-efficiency [104]. Similarly, [115] presented an improved avalanche beacon that lets rescuers know the precise position and the condition of buried avalanche victims.

3.2.3 Healthcare

Section 2.3 has already presented some of the many possible ubiquitous computing applications in healthcare and assistance for the disabled.

At home, a smart medicine cabinet [54] that is aware of its contents and of the current prescriptions for all members of the family, will remind us of the medicine to be taken, and issue warnings whenever a medicine is taken out that has not been prescribed. It would also know when a medicine has been recalled from the market, for example when new studies have shown dangerous side-effects unknown before. Going one step beyond (and having even more "self-consciousness"), a smart blister pack [56] would know when a pill has been taken out – as opposed to "only" know that a pack of medicine has been taken out of the cabinet.

Ubiquitous computing is likely to bring an increase in the everyday independence of several groups of patients. People suffering of dementia will enjoy a larger independence when smart environments [6, 116] will monitor some of their activities, reminding them of things they forgot, and issuing warnings to the paramedics when the sensors show a deteriorating condition. "Computer-Supported Coordinated Care" [24], an initiative started by Intel, has similar aims for elderly persons. Another group that is nowadays typically bound to the hospital, are patients suffering from a relatively high risk of a sudden fall in their otherwise stable condition – a decline that would quickly lead to a life-threatening situation. An early ubiquitous computing vision from the CoolTown project, to drastically improve the life of such persons by having on-body sensors monitoring their condition and alarming paramedics in case of an emergency, has been recently put into practice by a Swiss company.[4] Another company could even imbed all the monitoring hardware in a piece of garment.[5]

But not only sick people will take advantage of the healthcare benefits of ubiquitous computing. For healthy ones that want to stay that way there are also an entire range of possible applications. The above mentioned monitoring garments can be used not only for elderly or for observing outpatients in post-operative and chronic illness situations, but also as training support for athletes, for hazard materials workers, or for soldiers. They could even be used for watching professional truck drivers' vital signs to alert them of fatigue.

Finally, numerous ongoing projects use ubiquitous computing technologies to help the physically and mentally disabled. A wearable visual aid warns the blind of obstacles in their way [20]. Wearable navigation aids help the visually impaired to find their way through different environments, like cities or campuses [68, 123]. A large step towards the

---

[4] See www.auricall.com.

[5] See www.sensatex.com.

social integration of mentally disabled persons would be to smoothly integrate them into the games of others. Thereby, the other players don't have to play with "open cards," neither to play against a weak opponent, since the mentally disabled would be helped by electronic assistants [15, 55].

3.2.4 Staying in Touch

Staying in touch with loved ones over large physical distances, knowing about their condition, is another promise of ubiquitous computing. The wide-spreading of the Internet, together with e-mails, instant messaging, cheap or even free telephony, and webcams, has already shown how deep the societal impact of communication technology can be.

Through the use of ubiquitous computing, however, entirely new channels for experiencing closeness to family, friends, or lovers that are physically far away could be opened. Numerous prototypes have explored such possibilities over the last years. An early project, "LumiTouch" [22], has shown how with the use of relatively simple ubiquitous computing technology a large impact can be achieved. Since many people have pictures of their loved ones at home or on the office desktop, the researchers from the MIT Media Lab have used picture frames that can lighten up and are equipped with touch sensors. The frames are paired over the Internet. When one user touches the own picture frame, the other one will lighten up for some seconds, slowly fading after. Thus, the remote family member will know in a very unobtrusive way that the other thought of him or her, and can even respond by taking the same action.

Having had a slightly different focus, the CoolTown project envisioned picture frames that can emit light as well, however in different tones and intensities. They would typically be found in the homes of grown-up children that live far away from their elderly parents. The color and intensity of the parent's picture frames would – unobtrusively as well – provide information about the health condition of the parents, being connected to their in-home health monitoring system. The children can thus know about the parents' good health condition without the need of explicit questions, that might bother the parents and hinder other, more positive discussion topics.

A more recent project has brought along the so-called "Lover's Cups" [23]. The Lover's Cups are drinking glasses that have been enhanced with touch sensors on their sides and on the top surface that is touched by the lips when drinking. Thus, they know when they are being hold in the hand and when is being drunk out of them. The cups further exhibit accelerometers and temperature sensors, thus being aware of the moments when they are shaken and the contained fluid temperature. Finally, they can wirelessly get connected to the Internet (and thus, to each other), and they can display via LED's information from other cups. The envisioned scenarios for the Lover's Cups include: groups of friends toasting together over distances, friends or family members working in different offices taking a coffee break "together," lovers noticing when the other one drinks at a distance, possibly taking the same action in return (which would get noticed back), resulting thus in a more intimate connection between the two. The connection results from the knowledge of the other's actions and/or the shared action at one precise moment in time.

Common to all these examples is the fact that they provide a more unobtrusive, but at the same time more intimate channel of communication than classical telephony or instant messaging. These subtle forms of communication can further happen in parallel to other activities, in a business meeting for example, giving the participants the feeling of more continuous closeness and in some cases even the feeling of a shared secret that only they know about. Such new channels of communication over distances can obviously either be used on their own, or complementing nowadays forms of communication, such as voice or video calls. Ubiquitous computing promises thus to enrich the experience of close people that live at distance with several new dimensions.

3.3 Societal Risks of Ubiquitous Computing

The very means by which ubiquitous computing offers all the impressive possibilities presented throughout this chapter, however, could also lead to severe societal drawbacks. The possible perils of ubiquitous computing, that have increasingly concerned many, in both academia and society at large, are presented in this section.

3.3.1 Privacy

Loss of privacy is one obvious concern.[6] The generalized, 24/7, all domains of everyday life encompassing surveillance that can potentially be achieved by ubiquitous computing technologies is the enabler for many of the security, safety, and healthcare applications presented so far. It can, however, also lead to a state where many of the privacy borders that evolved in modern societies over centuries could be endangered [97]. Some even argue that a totalitarian claim stays at the very core of ubiquitous computing, since the vision of its pioneers is to explicitly intrude upon all aspects of everyday life with (surveillance) technology. Araya, for example, sees in ubiquitous computing "an attempt at a violent technological penetration of everyday life" [7]. "The old sayings that 'the walls have ears' and 'if these walls could talk' have become the disturbing reality. The world is filled with all-knowing, all-reporting things," argues Lucky [105]. And Adamowsky even states that ubiquitous computing is "a project that aims at totality and, of course, verges on the totalitarian" [2].

3.3.2 Health and Environment

Looking, among other societal risks and opportunities, most prominently at the consequences for human health and the environment, [75] warns about the possible negative impact of pervasive computing on individual health (due for example to the stress induced by the novel technologies, but possibly also by non-ionizing radiation, whose effects on the human body are not yet entirely understood), to the effects on the healthcare system due to the possibly increased costs, as well as the environment-related drawbacks. The latter could arise both due to the increased amount of electronic waste, but also as a consequence of the so-called "rebound effects." Rebound effects denote that a technology that decreases the need of one resource, usually makes that resource cheaper, which in turn leads to a higher demand for it, often increasing its consumption in the end rather than decreasing it. For example, the already mentioned "intelligent" navigation system that navigates vehicles around traffic jams, on a first level obviously results (on average) in less fuel consumption for one journey. However, due to the also im-

---

[6] And, as such, by far the most thoroughly regarded upon in literature, e.g. [65, 95, 96, 129]. It will thus become the object of this thesis only at one point in chapter 6; it is however mentioned here due to its outstanding importance and in order to provide the broader picture.

plied time economy, drivers would then tend to plan longer (or more frequent) journeys, since they know they will now take a shorter time. Finally, this could result in more fuel consumption than before [75].

3.3.3 Social Compatibility Issues

Bohn et al. [17] identify an entire set of potential societal problems, divided into three sets: reliability, control, and social integration issues.

Reliability

Instead of one bulky device containing a large number of features, ubiquitous computing environments typically contain a multitude of small, individual devices, from a multitude of manufacturers, based on heterogenous technological platforms, that work together to provide the required functionality. However, this can result in a multitude of reliability-related problems:

- *Manageability*: Implementing, controlling, managing large-scale ubiquitous computing scenarios involving potentially millions of smart, adaptive devices, will be a large challenge for programmers. As [31] put it, "What happens when driver 3.06.2 proves to be faulty and has to be updated with version 3.06.3 on all 15 billions temperature sensors in this world?" On the other hand, understanding the way such systems work will most likely be a huge challenge for users.

- *Predictability*: An important feature in using tools is their predictability. Ramming an axe in a piece of wood is a very predictable action, with two possible outcomes: success (the piece of wood has been cut in two) or not (the axe remained stuck in the wood). For both successful "operation" of the axe and failure, the result is obvious. For technology tools, where such predictability is not naturally given, it is important for the designers of the system to provide the users with simple means to learn and predict the systems. For a telephone, for example, lifting the receiver is the straightforward action to start a call; when no dial-tone comes, the system is broken. Due to the unobtrusive nature of ubiquitous computing systems (that are envisioned to quietly work in the background), however, it might be difficult to provide users with means of ascertaining their state.

Control

The explicit goal of ubiquitous computing to hide the actions from users, and only provide them with the service, not only complicates the reliability issues presented above. It can also have more subtle societal implications:

- *Content Control*: Many envisioned ubiquitous computing applications rely on smart objects providing information about themselves to other smart devices, and finally to the user. Due to the unobtrusive, hidden nature of those interactions, however, it will ultimately be almost impossible for the user to assess the accuracy of the received information. Bohn et al. [17] identify several examples of possible misuses: "For example, smart products might be used to tie customers more closely to traders by recommending they purchase other goods produced by the same trader." And further, "maybe more importantly: who will decide what a smart toy tells the children, potentially shaping the children's opinions without their parents' knowledge? Tempting children to buy additional toys would only be the most obvious strategy – a much more serious threat would be the moral values induced by smart toys during play."
- *System Control*: Due to the same unobtrusiveness, it might finally be possible that smart objects be not totally loyal to their owners anymore. As today's digital rights management systems control the way we handle media such as music or movies, future smart objects might limit not only the way we use our digital assets, but also our real-world belongings. Smart cars, for example, might control the way we drive by limiting their speed to the current speed limit, or not letting the doors open in a no-parking area. Which, when in need to quickly accelerate and exceed the speed limit to avoid an imminent accident, or to park in the no-parking zone in front of a hospital to urgently drop the pregnant wife who is about to give birth, may come in unhandy.

Social Integration

Another fundamental challenge for pervasive and ubiquitous computing systems is their social compatibility, that encompasses issues of system transparency, social equity, and acceptance. Examples for such issues, as identified by [17], are:

- *Transparency*: Many of the economic drivers for ubiquitous computing presented in section 3.1 rely on smart artifacts autonomously triggering actions, that in practice could often involve costs they would have to pay in our name. Detecting malfunctions of such paying systems, but also keeping track of all the payments executed by our many smart devices could prove as tough challenges.
- *Knowledge Sustainability*: The more dynamic the world around us gets, the less useful past experiences (e.g., food prices) become in dealing with the present. Acquired knowledge becoming obsolete at an ever-increasing pace could, in the long term, contribute to an increased uncertainty and lack of direction for people in society.
- *Fairness*: The personalization of services enabled by future ubiquitous computing environments not only has bright sides. Some customers might not be worthy to be offered certain services or premium information. David Lyon, Professor of Sociology, calls this process “social sorting” – “Categorizing persons and groups in ways that appear to be accurate and scientific, but which in many ways accentuate differences and reinforce existing inequalities” [107].
- *Universal Access*: Finally, the increased complexity of systems and the sheer amount of devices in future could counter-balance the expected easiness of access to information, ultimately increasing the digital divide separating groups like elderly, or less technologically inclined from the rest of society.

3.4 Summary

This chapter made the case that on top of the technological drivers presented in chapter 2, there are also strong economic and social drivers for the deployment of ubiquitous computing (i.e., business interests and the potential of important societal benefits, respectively). As shown later in the chapter, however, these could be counter-balanced and even outweighed by a multitude of different pitfalls that could come along with the deployment of ubiquitous computing.

However, so far only a broad overview of both risks and benefits has been presented. How their specific incidence probabilities and magnitudes can be realistically assessed (so that the most relevant are regarded with priority), what measures can be taken to counter the most alarming threats, which are the relationships and tradeoffs between technological and societal challenges – all these questions that stay at the core of assessment of ubiquitous computing implications have been let aside. Finding answers for them will form the remainder of the thesis.

4 Assessing Implications of Ubiquitous Computing – Existing Methods

Lessig's statement that "code is law" [101] points out the societal responsibility of IT designers and developers. Assessing the societal implications of IT systems has thus been an important part of the broader field of technology assessment for quite some time now. However, the expected large-scale fielding of ubiquitous computing applications will presumably – as argued in the previous chapters – have significantly more far-reaching societal consequences than the IT systems used today. Smart things and smart environments will probably govern many aspects of our everyday lives on a 24/7 basis, and not only our online presence. They will often become indistinguishable from the environment, so that people could be taken by surprise that an object has displayed some smart behavior when not expected to do so. The repercussions of such an extensive integration of computer technology into our everyday lives are difficult to predict. Trying to make precisely these predictions, however, is imperative. As Hilty argues, "the reason for assessing technological risks at an early stage of development is the experience that full evidence for the existence of a risk is sometimes only available after severe damage has occurred. For example, the fact that chlorinated fluorocarbons (CFCs) harm the atmospheric ozone layer was not proven until after the ozone hole had been discovered" [77].

While some effort has been undertaken within the community to analyze the possible privacy threats of ubiquitous and pervasive computing systems, many other societally relevant issues remain largely unaddressed. Moreover, within such a multitude of opportunities and dangers as presented in the previous chapter (and among the many more that could not be presented), it truly is difficult not to lose orientation. Consequences of ubiquitous computing are far-reaching, various and heterogenous, and difficult to predict. Which of the ubiquitous computing applications that would bring along these risks and opportunities are likely to become reality, which are not? Even if such applications will arise, how large is their potential to change societies, i.e., how large are their incidence probabilities and magnitudes? Which threats only arise through the combination of several ubiquitous computing technologies that emerged independently for different purposes? How can the most alarming dangers be fought against: are technological measures possible and do they suffice? Must laws be updated to enforce some technologies and channel others? Or is it not technologies, but the processes (e.g., economical processes) that have to be legislated? Are solutions needed that combine technological and political measures? How large will the government's greed to take advantage of the surveillance be and how can it be controlled by society?

For answering these and similar questions, we will propose in the next chapter the method of *Scenario-Driven Prototyping*. Before that, however, we use the current chapter to present the approaches used so far to understand the repercussions of ubiquitous and pervasive computing. Besides future scenarios – a tool widely used for the implication assessment of various technologies – for the specific case of ubiquitous computing many researchers have used a procedure which we call "technology-driven prototyping." The chapter starts by presenting this latter method in section 4.1. Section 4.2 then presents an overview of the history and methodological development of future scenarios. Section 4.3 narrows the discussion to ubiquitous computing related scenarios only. Finally, section 4.4 summarizes the presented approaches, highlighting their merits and weaknesses.

4.1 Technology-Driven Prototyping

Ubiquitous computing is a fairly novel and rapidly evolving field. Since it also encompasses a large number of technologies, it is often difficult for persons from outside the field to keep track with its developments, or to foresee their implications. However, precisely due to this high dynamism, researchers from within the field are often bound to show the possibilities of the novel technology by creating prototypes of some possible future applications. Such prototypes sometimes also allow an analysis of the societal opportunities and challenges they encompass.

This path, that we call *technology-driven prototyping*, is thus typically taken by researchers with technical background: from a collection of novel technologies, they build a prototype that stands as showcase for some possible future application making use of them. Typically, this approach is more or less well-suited to identify the technological issues that would arise by extending the prototype to a large-scale, industrial-strength application.

Sometimes, societal implications of the considered application also appear as a side effect of looking at the possible use cases of the application. However, more often than not, the analysis is very shallow, many other concerns being more urgent for the engineers developing the prototypes. Langheinrich [97], who visited seven large European ubiquitous computing projects to analyze their specific efforts towards privacy, noted: "As long as privacy is situated on a non-critical development path, more important issues such as energy efficiency, code size, or robustness dominate the researcher's todo-lists. Decisions pertaining to data storage and communication details are often improvised and seen as a temporary solution fit for prototype deployment. Projects which explicitly had privacy issues as part of their deliverables, generally exhibited greater concern for such issues, even though they often stopped short of generating novel ideas and limited themselves to a broad but shallow summary of general privacy issues, without taking project specific design parameters into account." And further: "Most researchers who participated in the interviews and discussions did not (yet) think of privacy issues in their own work, or only on a very obvious level." [97].

However, it is not only constraints of time and resources that make the analyzes of societal implications resulting from technology-driven prototyping to be often irrelevant. Even when the engineers take the societal consequences of their prototypes seriously and analyze them profoundly, there is a more subtle effect that can happen. Very often, such prototypes emerge from a purely technological feasibility background, without any concern about basic societal compatibility questions (e.g., would we really want to rely so much on technology and lose all common sense, so that when technology fails to work we would be lost, as some projects suggest?), or concerns for the economic feasibility of the presented application (i.e., is there a realistic business model behind the application? Would any company invest to develop such application so that we need to be concerned with its implications?)

It is then precisely these unanswered questions that Davies and Gellersen identify as being one of the main causes that, despite the tremendous amount of technological progress over the last decade, "many aspects of Mark Weiser's vision of ubiquitous computing appear as futuristic today as they appeared in 1991" [39]. Altogether, the authors make out three major factors responsible for most of Mark Weiser's visions not yet having become reality. One is a technical limitation, the fact that many prototypes developed within the field require the system to know what the user intends to do – which is also an old unfulfilled dream of Artificial Intelligence. The other two causes are of societal nature: the unanswered societal questions, and the poorly analyzed and often naive economics and business models that form the basis of many of the prototypes, making many of them seem more of economic utopias, despite being technically feasible.

Lastly, numerous prototypes developed within the field don't even get to the point where they could be regarded as serious candidates for answering societal issues of ubiquitous computing. They have in fact a narrow scope, addressing only the world directly known to their developers and are thus typically concerned with small convenience improvements in the office life of the typical white-collar worker. Hence, they are ill suited to assess any broader societal opportunities or possible pitfalls of ubiquitous computing, nor can they provide much help in analyzing what technical approaches are likely to facilitate the former and discourage the latter [30].

4.2 Future Scenarios

By contrast, the *future-scenarios-based analysis* is the path typically taken by social scientists who intend to analyze the societal implications of upcoming technologies from their specific scientific viewpoint. The current section presents an overview of the historical development and the methodology of future scenarios.

While "future studies," as a systematic and scientific[1] approach for picturing possible futures, has emerged relatively recently, its roots lie in the much older utopian future literature, whose roots Bell [12] traces back to the 16th century.

We will thus start this section by presenting some of these utopias – however, only technology-related utopias from the last one hundred

---

[1] Whether the term "scientific" should be used when, as Blass points out "the very notion of researching the future is a paradox. The word research lies within the time boundaries of the past and the present, so to research the future appears a logical impossibility," [14] is a debatable matter, outside the scope of this thesis. For a discussion, see the same article of Blass [14]. Nevertheless, as this section will show, many serious academic and industrial institutions have been concerned with future studies over the last decades.

years. Then, we will proceed to the scientific and systematic approaches to envision future worlds and implications of technology, as they have emerged in the second half of the 20th century, and present a brief history. The section ends by focusing on the few well-known, recent, ubiquitous computing-related scenarios, which are closely related to the topic of this thesis.

4.2.1 Utopian Visions of the Future

While, as argued above, future utopias have been around for some centuries already, the more technology-oriented ones (and thus of more interest within the context of this thesis) have started after the industrial revolution in the 19th century.

Arguably, the father of modern technology-related visions of the future has been the French novelist Jules Verne. In one of his most well-known novels, for example, “From the Earth to the Moon” [159], he envisioned Man’s first travel to the moon hundred years ahead of time. While the method used to enable the “lunar module” to reach the escape velocity from Earth’s gravitational field – a shot from an immense gun implying enormous accelerations – is rather questionable from an astronaut’s perspective, Verne’s vision is scientifically founded, so that quite an astonishing number of parameters (in historical retrospective) are similar to the Apollo 11 mission: in Jules Verne’s novel, as in later reality, the launch site is located in Florida, the southern-most point of continental US (and thus with the largest Earth’s rotational velocity), and a crew of three embarks for a one-week long trip. Moreover, the dimensions of Verne’s man-sized “bullet” are very similar to the ones of the lunar module that would ultimately land on the Moon in 1969.

Apart from Verne’s early utopias, two periods in the recent history – both characterized by a period of peace, fast technological progress, and economic prosperity within the western world – have as a result animated numerous authors to write optimistic utopias about the future, a future in which constant technological advances not only make life more pleasurable and prosper, but also help solving various societal problems.

Fin-de-Siècle Technology-Utopias

The first such period has been around the end of the 19th and the beginning of the 20th century. In a climate of rapid technological and scientific advances during the industrial revolution, nations used to display these progresses at the then-popular world fairs [113]. Well-known examples are the "Expositions universelles" in Paris from 1889 and 2000 – as a display of technological achievement by the French, for example, the fair of 1889 brought Paris its landmark, the Eiffel tower. Scientific discoveries and technological developments from that period include: the discovery of radiation by Pierre and Marie Curie, of electricity by Graham Bell, and further such technologies as telegraphy, gramophone, telephone, or cinema.

Based on such advances, it became popular at the dawn of the 19th and start of the 20th century to write utopian, technology-based visions of the world as it would be in one hundred years from then, in the year 2000. The authors of such visions usually envisioned a happy future, in which manual labor has almost disappeared, machines and robots taking over most of the tasks, the technology allowing for a better, happier, more harmonious life, without poverty or diseases.

Mattern presents in [113] numerous examples from the 100-year-old utopian literature, which envisioned among others future ubiquitous personal flying machines, cleaning robots taking over all household duties, or the "cinema-phono-telescope," a combination of video-telephone and virtual-reality-device, allowing people to communicate over distance in the same way as if they were together in one room, thus rendering travelling superfluous. Several authors from that period envision the ubiquity of such "wireless telegraphy," and see it's implications as exclusively positive for mankind. One such author for example, writes: "Monarchs, chancellors, diplomats, bankers, and directors will be able to attend their businesses regardless of where they are. Directors of a company will hold a meeting, one of them being on the top of the Himalayas, the second at the beach, and the third during an air travel. They will see each other, talk to each other, exchange their documents, and sign, as if they would be together in one place." And further: "Nowhere, regardless of the place, one will be alone any more [...] In future, the wife will be able to assess what the husband is doing at any moment; and he will himself be able to see whether she only thinks about him. Lovers and married couples will never be apart, even if hundreds or thousands of miles lie between them. They will always see and talk to the other – in short, an era of happiness for love will emerge." ([144], as cited by Mattern [113])

Prosperity and Technological Progress after WW2

The second period of technology-based future optimism has been the time immediately following World War 2, a time characterized again by stability, material prosperity of the western world, and a wide-spreading of novel technologies into society and people's homes. The 1950s and 1960s brought for example the civil use of nuclear power, the rapid spreading of motoring and of affordable air travel, or the fast-paced conquer of the outer space and the Moon (made possible by huge investments).

The future predictions from this period are remarkably similar to the ones that had been made half a century earlier. The year 2000 would see solar-powered flying vehicles (or ones powered by small nuclear engines), household robots would take care of the house duties, some of us would be living in underwater cities or in lunar colonies, while the ones still on the surface of the planet would work 10 hours per week, travel with the small family helicopters, and spend weekends on the Moon [78]. Cancer would have been eradicated, as would have been all infectious diseases [151], while no new infectious diseases like AIDS, BSE, or bird flu have been foreseen.

On the other hand, almost all of these predictions have underestimated the role that computers and communications technology would finally have on our present society. Many were seeing the computers of the year 2000 as still working with punch cards, and even if some predicted the ubiquity of mobile phones, they have – as Mattern analyzes – always been seen as means to attend businesses or solve other relevant issues: "That twelve-year old school girls will send each other text messages over two meters of distance, or even snapshots made by their camera-equipped mobile phones, was simply not conceivable at the time" [113]. While most progresses have been overrated, the "digital era," with mobile phones, cameras, PDAs, mp3-players, blogs and chatrooms has not been predicted.

Technology as Universal Problem Solver

Besides some sporadic technological naiveness (or the difficulty to foresee disruptive technologies such as the World Wide Web), it seems in retrospective that the predictions from both these periods have been so wrong because of the same two naive assumptions. The assumptions relate to the technological optimism of the predictions and the fact that they have ignored any factors other than the technological advances themselves. The technologically possible would be done and would equally benefit everyone in the society.

Firstly, technology is seen as a sort of a universal problem solver, which only benefits societies, and does not bring negative effects. To cite Mattern again, "environmental pollution, overpopulation, unemployment, pandemics, hunger, poverty, terrorism, surveillance society, globalization or climate change are obviously no issues for the year 2000 [113]." Technology would solve most societal and health problems, without inflicting new ones. The old infections will be cured, no new ones will appear. Being able to reach (and thus control) one's partner at any time, is genuinely positive and implies no negative effects. Societal developments other than towards general happiness were also not foreseen: the wife would still happily and smilingly wait home for the husband to come home from work, just with a more pleasant daily life, since roboters would take care of all household duties [31].

Secondly, economic viability is no issue. The fact that the race of the two superpowers to the Moon had only been possible by enormous expenses is ignored in such predictions; and thus the fast technological advances from the 1960s are extrapolated with the same (or higher) pace for the future. That a ten-hour work week is unlikely to guarantee a sustainable economic growth (or, indeed, the scientific and technologic progress itself) does not seem to bother. Neither that in course of the globalization, countries with cheap workforce could conquer some of the markets, yielding many of the fathers without work to come home from to their smilingly waiting families.

4.2.2 RAND and Herman Kahn – First Systematic Approaches

The first attempts at foreseeing likely futures based on a systematic approach, have their origins in the military strategies developed during World War Two and shortly after. As Kleiner notices, "World War Two has been an extraordinary catalyst for the study of complex systems. Just as the war had mingled social scientists in unprecedented numbers, it had also assembled physicists, mathematicians, logicians, and psychologists to work on problems beyond the range of any individual discipline" [89]. These scientists started to apply game theory and decision analysis techniques for interdisciplinarily answering military questions such as how to patrol a stretch of water with the minimum amount of boats, or what the enemy's likely responses to different courses of action would be [89]. Especially this second type of questions represents a first restricted attempt at scientifically understanding how the future could evolve. Gradually though, from these first narrow questions, the technique of scenarios as a tool to assess possible futures would evolve, and it will seek answers to increasingly complex questions.

Herman Kahn: Thinking the Unthinkable

Some of the scientists working with game theory during the second World War joined shortly afterwards the RAND corporation,[2] newly founded in 1946. Inside this think tank, they refined their methods for answering the new strategic military questions posed by the cold war.

Most known among the young scientists working at RAND has become Herman Kahn. Kahn used game theory to speculate about the possible series of events during and outcomes of a thermonuclear war, arguing that such a war is not only survivable by human civilization, but also winnable, and that an outcome with a few major US cities destroyed is better than an outcome with all major cities wiped out. In other words, he was "thinking the unthinkable," as Scientific American wrote at the time [63]. Kahn, who seemingly loved to provoke as much as he was convinced that the best way to avoid a nuclear war was to deeply and dispassionately think about its consequences [63], gladly adopted this term and published his ideas 1950 in the book "Thinking About the Unthinkable" [84]. Ten years later, he published the same ideas in a considerably larger version, with the straightforward title "On Thermonuclear War" [85].

The RAND Corporation and the Delphi Technique

However, in 1960, when "On Thermonuclear War" appeared, Herman Kahn and the RAND corporation had already started to move on, widening both the objects and the methods of their studies. While still looking into the future, they were no longer interested in military questions only, but in more general societal questions. They also did not further rely solely on their own knowledge of the fields studied, but proposed the so-called "Delphi" technique, in which various experts

---

[2] See www.rand.org.

are interviewed about the time horizon in which they expect certain technological or societal developments to happen.

The first-ever Delphi study [64] has been published by Gordon and Helmer-Hirschberg from the RAND corporation in 1964. It treats six broad technological and societal areas: scientific breakthroughs, population growth, automation, space progress, probability and prevention of war, and future weapon systems – the last two issues being obvious tributes to the military and strategic origins of the RAND corporation. The authors explain in the beginning of the study how this expansion of interests came naturally, and what the scope of the study is: "Substantively, our interests lay in assessing the direction of long-range trends, with special emphasis on science and technology, and their probable effects on our society and our world. Here, by 'long-range,' we had in mind something of the order from ten to fifty years. Our natural curiosity in this regard was enhanced by an awareness of the fact that our work at RAND is in many instances closely related to plans and policies affecting the rather distant future" [64].

They also acknowledge from the outset two limitations inherent to such studies. The first one is the obvious uncertainty of the future: "In seeking out the future trends in these areas, we were of course well aware that we would not through some miracle be able to remove the veil of uncertainty from the future. This did not seem to us to imply, though, that it is altogether impossible to make meaningful assertions of substantive content about the future." The second is the more specific unpredictability of revolutionary evolvements or breakthroughs: "Future events can be considered as roughly belonging to one of two sets: the expected and the unexpected. A study such as this cannot hope to uncover unexpected, spectacular, unanticipated breakthroughs, but must concentrate on narrowing down the dates and circumstances of occurrences which can be extrapolated from the present" [64].

Since the scope of the study is very broad, while we are more interested in the technology-related future studies, we will not present here in detail its results. Nor will the exact Delphi methodology (consisting of several rounds of interviews between which the experts are anonymously presented the arguments brought by others so that they can ammend their own predictions in light of the different opinions) be presented. Important about the latter is that the study offers *quantitative* results, i.e., it provides a time interval in which with high probability (according to the experts), a certain technology will have become available or a certain societal event will occur. As for the results of this first Delphi study, we only provide one example here to show the relevance of the problems discussed and the foresight of some results. Thus, the study foresees "the eventual abundance of resources of energy, food, and raw materials, but also the possibility that a continuing inequitable world distribution of these assets to the increasing world population may furnish a persisting stimulant to warfare" [64].

A few years later, in 1967, Kahn and Wiener published the book "The Year 2000: A Framework for Speculation on the Next Thirty-Three Years" [69], based on a Delphi study as well, and focusing specifically on the technological novelties likely to emerge in the then-remaining last third of the 20th century. As Mattern analyzes in [113], most of their predictions have been far too optimistic. He lists some examples from the top one hundred advances expected by the authors: new energy sources for road and train transportation (such as magnetic levitation), extensive use of cyborg methods, the use of atomic blasts in mining, human hibernation for extended periods of time, synthetic aliments, manned moon outposts and interplanetary travel, underwater colonies, extensive use of robots as workforce, personal flying platforms, and novel methods for the easy learning of foreign languages.

As we will argue in more detail later, one of the likely reasons for these erroneous estimations is the exclusive focus of the study on the technologically possible. The fact that the technological progresses during World War Two (leading to the quick development of such technologies as the radar or the atomic bomb), or during the Cold War (leading to fast-paced progresses in space technologies) were only possible due to a virtually unlimited funding and concentration of scientists in these areas, and that such concentration is not to be expected in a different societal setting, has most likely not been considered. After such a period of one-sided efforts of the society towards technological advances, it is more likely than not that societies will have different priorities, such as economic prosperity, education, fighting the poverty, or caring for the environment. And even within research itself, as we have argued in [31], the society is likely to invest in the areas where it sees the largest opportunities for societal welfare. While the potential benefits of – for example – medical, genetical, biological advances are rather obvious, the deeper societal meaning of moon or underwater colonies is more questionable. Finally, the fact that the experts invited by Kahn and Wiener were exclusively scientists from engineering disciplines, together with some science-fiction authors [113], did probably not help in bringing this needed larger economic and societal perspective into their study.

Nevertheless, these first Delphi studies of the RAND corporation are widely regarded as the birth hour of scenario planning (e.g., by [89, 128, 140]). They have been the first systematic and scientifically-based approaches to understand future consequences of today's actions and efforts, while at the same time addressing the inherent difficulties of doing so. They have also understood that, in order to be a helpful tool for decision-makers, several possible futures have to be presented, together with the individual decisions that lead to any of these possible futures. Thus, as opposed to the early war simulations done by Herman Kahn, the means to present such possible futures have became more narrative, as a collection of stories, emphasizing which courses of events could lead to one or the other story. Also during this period, the term "scenario" has been coined for such future stories. Based on Hollywood terminology, the name has been suggested by the screenwriter Leo Rosten [90], and, as Ringland observes, "Herman Kahn adopted the term because he liked the emphasis it gave, not so much on forecasting, but on creating a story or myth" [128].[3]

4.2.3 The Shell Story – Success through Scenarios

Telling myths through stories about the future is also what the multinational Royal Dutch/Shell oil company (called simply "Shell" from now on) started to do in the early 1970s. In an increasingly complex and instable world, the planners at Shell recognized that straight business forecasting, as had been done for the 25 years before, would no longer suffice for the strategic planning of the company. Inspired by Herman Kahn's work, they introduced the method of scenario planning, which lead – as the current section will show – to quite some success stories.

While after Shell's successes, several other companies started to use future scenarios to asses their future business environments, as well as the according opportunities and risks, the story of the Shell scenarios is still the most interesting one. Not only since by using them Shell has

---

[3] As a matter of fact, as Kleiner describes, Herman Kahn was friends with quite some directors and screenwriters from Hollywood, the RAND corporation being located in Southern California as well. One of them, Stanley Kubrick, is credited to have taken him as inspiration for the fictional character "Dr. Strangelove," from his movie with the same title [90]. As the movie character, Kahn had thought of a doomsday machine as the ultimate way of discouraging the other superpower to make a first nuclear strike [63].

succeeded to foresee and thus to be prepared for some major geopolitical shifts with heavy influence on the worldwide oil business, but also because their scenarios had far-reaching consequences beyond the limited world of business. On several occasions, the Shell scenarios have successfully influenced entire societal developments into a positive direction, such as peaceful development of a troubled area. They thus came close to what Roßnagel sees as the ideal of future scenarios: "provide a rational answer to the question: Which future shall we want" [134].

First Scenarios at Shell – Foreseeing the Oil Crisis

The story of Shell's future scenarios is closely related to the name of the frenchmen Pierre Wack, head of Shell's planning department throughout the 1970s, who, together with his colleague Edward Newland, introduced the scenario technique as strategic planning method for the company. Wack described his experiences in a series of two articles published in the Harvard Business Review: "Scenarios: uncharted waters ahead" [161], and "Scenarios: shooting the rapids" [160].

After describing how the straight forecasting method ("explore and drill, build refineries, order tankers, and expand markets" [161]) had been used by Shell during the relatively stable 1950s and 1960s, Wack, who was familiar with Herman Kahn's work and intrigued by its possibilities for corporate planning, explains how they ended using it within Shell. During the late 1960s, while still working for the French planning department of Shell, Wack and his team were confronted with two major uncertainties of the French market: the availability of natural gas, with potential large side-effects on the oil business, and the future of French politics. France used to be protectionist to the advantage of national oil companies – and limiting foreign companies', such as Shell's, market share. In the years to come, however, as a member of the European Community (EC), France might have had to change its policy and liberalize the market to the benefit of other European companies. Obviously, under such uncertainties, straight forecasting could not longer work, and Wack was allowed to use scenarios to depict the possibilities. Combining these two uncertainties (high vs. low availability of gas, and protectionism vs. liberalization) into a matrix of four possibilities, Wack and his team presented the Shell management with a set of what he later called "first-generation scenarios." While they clearly presented the four possibilities and what their specific consequences would be, they were not a good tool for decision makers: "More important, we realized that simply combining obvious uncertainties did not help much with decision making. That exercise brought us only to a set of obvious, simplistic, and conflicting strategic solutions [...] In emphasizing only uncertainties, and obvious ones at that, the scenarios we had developed were merely first-generation scenarios. They were useful in gaining a better understanding of the situation in order to ask better questions and develop better second-generation scenarios – that is, decision scenarios" [161].

For developing second-generation scenarios as a better tool for decision-making, Wack realized the importance of discovering and then highlighting within the scenarios not only the uncertainties, but also what he calls the "predetermined elements" of the future – that is, the certainties. He writes: "Yet this negative realization led to the discovery of a positive search tool. By carefully studying some uncertainties, we gained a deeper understanding of their interplay, which, paradoxically, led us to learn what was certain and inevitable and what was not." And further, recognizing the limitations of the scenario technique, as Kahn had done before, but further refining it by limiting the possible futures through the careful recognition of predetermined elements: "Strictly speaking, you can forecast the future only when all of its elements are predetermined. By predetermined elements, I mean those events that have already occurred (or that almost certainly will occur) but whose consequences have not yet unfolded. Suppose, for example, heavy monsoon rains hit the upper part of the Ganges River basin. With little doubt you know that something extraordinary will happen within two days at Rishikesh at the foothills of the Himalayas; in Allahabad, three or four days later; and at Benares, two days after that. You derive that knowledge not from gazing into a crystal ball but from simply recognizing the future implications of a rainfall that has already occurred. Identifying predetermined elements is fundamental to serious planning" [161].

The concept of "predetermined elements," which complements the uncertainties, and has been used in the second-generation scenarios, has then led to the largest success of Wack's department – to foresee that an oil crisis in the 1970s is bound to happen – not as a possibility, but as a predetermined element. Wack, from 1971 leading Shell's global planning department, had as first task to look for many years ahead into the relation between oil exporting and oil importing countries. By considering different uncertainties and the individual scenarios resulting from these, they discovered two common patterns to all of them: while the economies of the oil importing countries would continue to grow at a steady pace (in some cases, such as Japan, even increasingly fast), with a need for oil growing faster than their growth in GNP (gross national product), the oil exporting countries would shortly either be unable or unwilling to increase their exports. Especially this second insight was a striking novelty. While, as Wack analyzes, it was clear for everyone in the industry that some countries, such as Indonesia, could not further increase their production, it had been assumed that the exporting countries that could increase their exports would do so, following the increased demand, as they had done in the past.

However, the scenario method gave Wack's team the opportunity to take into account the interests of the various stakeholders of the oil business, among them the national interests of all exporting countries. By looking at two criteria (available oil reserves and absorptive capacity for the money resulting from sales), they saw that there was not one single oil exporting country that had both the reserves needed to face the increased demand and the capacity to absorb the revenue from sales. As Wack writes about Saudi Arabia, for example, "production would generate more revenue than the government could purposefully spend. We concluded that even though oil company logic would have the Saudis producing 20 million barrels per day by 1985, the government could not do so in good political conscience. It was no surprise when Sheikh Zhaki Ahmed Yamani, Saudi Arabia's minister for oil affairs, later remarked: 'We should find that leaving our crude in the ground is by far more profitable than depositing our money in the banks, particularly if we take into account the periodic devaluation of many of the currencies. This reassessment would lead us to adopt a production program that ensures that we get revenues which are only adequate for our real needs' " [161].

By the end of 1972, one year before the oil crisis became reality, Wack was thus able to inform the Shell management that shortly, within the next few years, a market imbalance was bound to happen – not as a possibility, but as a *predetermined* element. The continuously increased need of oil on one side would meet the incapacity or the unwillingness to increase production on the other side, the oil market would change from a buyer's to a seller's market as result. "We did not know how soon it would occur, how high the price increase would be, and how the various players would react. But we knew it would happen. Shell was like a canoeist who hears white water around the bend and must prepare to negotiate the rapids" [161].

Wack's Refinements of the Scenario Technique

Over the years, Wack and his team have refined the scenario technique in several ways. Firstly, as discussed, they have recognized the importance of spending energy to recognize predetermined elements, as a mean to reduce complexity and increase the accuracy of scenarios.

Secondly, they have also changed the way to deal with the remaining uncertainties. The first-generation scenarios were a mere extension of forecasting, except that a single forecast was replaced by a probabilistic assessment of alternative futures. As Heijden [155] observes, this has not proven a fundamental advance over other forecasting approaches. The real advancement in the second-generation decision scenarios, aside from a better examination of the predetermined elements, lies in the way the uncertainties are dealt with: "The aim is to develop projects that are likely to have positive returns under any of the scenarios [...] Decisions are not based on one scenario being more likely than another; project developers optimize simultaneously against a number of different futures which are all considered equally plausible, and treated with equal weight" [155]. Heijden refers to business projects, but the concept can be extended to any other type of projects, such as societal. And talking about Wack's paradigm shift: "The first objective of scenario-based planning became the generation of projects and decisions that are more robust under a variety of alternative futures" [155]. Wack himself writes about the decision scenarios: "They are not a group of quasi forecasts, one of which may be right. Decision scenarios describe different worlds, not just different outcomes in the same world" [161].

Thirdly, in order to better understand both uncertainties and predetermined elements, Wack insisted on looking at the actors of the play, stakeholders as we would call them today, as well as their interests. To illustrate, again Heijden's words about Pierre Wack when faced with the question of oil availability: "Shell's technical people had concluded that supply availability was predetermined, the resource in the ground was plentiful, and the necessary number of wells could be drilled. But Pierre Wack was not satisfied with that answer. He looked behind it, considering the people who have control over the reserves who would be making the actual production decisions. [...] It was one of Pierre's great contributions to the scenario process that he insisted on looking at the people behind the decisions" [155].

Fourthly, being bound to present their results to the companies' management and thereby challenge fundamental assumptions of their managers (such as the one that a business environment will remain as it had been for over two decades), they have also learned a great deal about how to present such challenging new world views. Presenting in 1972 their results about how the oil business will change, they have decided to present a collection of six decision scenarios. Three of them (the so-called "A-scenarios") were predicting the change they foresaw, being different in details, such as how the governments of the individual oil importing nations will react. The three "B scenarios" were attempts to perpetuate "business as usual," offering "solutions" for how the coming unbalance of offer and demand would be avoided. To match the upcoming shortage of oil availability, the B1 scenario assumed a very slow economic growth in the oil importing nations, much less than the most pessimistic of predictions. The B2 scenario was an artificial construct, where "everything would somehow solve itself." B3 was called the "three miracles scenario": against all odds, enough oil fields would be discovered in a short time, "all major producing countries would happily deplete their resources at the will of the consumer" [161], and the usual short-termed increases in demand or shortages of supply would not occur any more. Since all of these "solutions" were at least unrealistic, if not absurd, Wack reports that by including the three B scenarios they succeeded much better into challenging the mindset of their managers than if they would only have included the change scenarios. As fourth insight from the Shell experience, Wack thus argues that there should always be scenarios evolving into opposed directions, either because an uncertainty requires them, or to better underline a predetermined element.

Fifthly, from experiences gained later on, Wack argues that while there has to be a number of different decision scenarios, six were too many: "Decision scenarios acknowledge uncertainty and aim at structuring and understanding it – but not by merely crisscrossing variables and producing dozens or hundreds of outcomes. Instead, they create a few alternative and internally consistent pathways into the future. Never more than four (or it becomes unmanageable for most decision makers), the ideal number is one plus two; that is, first the surprise free view (showing explicitly why and where it is fragile) and then two other worlds or different ways of seeing the world that focus on the critical uncertainties" [160].

Impact of Shell Scenarios

Since the Shell management had been warned about the crisis that was about to happen with less than one year in advance (a short time for such a then counter-intuitive insight to infiltrate into the culture of a large multinational company), it could not prepare early enough to avoid all of its negative effects. However, due to the scenarios, they were able to recognize the first signs when the crisis started to develop, and to answer quickly. As Heijden writes about the the Shell management, "they interpreted persistent signals from that part of the world as an indication of the unfolding of the crisis scenario, and so they made a number of critical strategic decisions. The most important decision was a change in refining investment policies to allow for the possibility that the crisis scenario was in fact playing out. While most of the refining industry needed years to decide that something really fundamental had happened, Shell moved immediately, switching investments well ahead of their competitors" [155].

After Wack left Shell in 1981, Peter Schwartz led the scenario planning department of Shell for a few years, before he went on to found the Global Business Network,[4] the most-known scenario-based consultancy nowadays. During his time at Shell, Schwartz was involved in some surprisingly accurate forecasts, too, which he describes in his book "The Art of the Long View: Planning for the Future in an Uncertain World" [140]. Due to the scenarios developed by Schwartz and his team, Shell has again been prepared for a fundamental change to come ahead of its competitors. While everyone in the business, after the oil shock from the 1970s, was overinvesting in new fields, Shell realized that in the 1980s a decrease in the oil price was likely to occur. Davis-Floyd writes: "In the early 1980s, one of the scenarios written by the Shell planners foresaw the likelihood of a rapid and dramatic decrease in the price of oil as the result of the discoveries of new fields outside of the OPEC sphere of influence, in combination with the energy conservation measures increasingly taken by consumers who did not want, after the debacle of the 1970s, to remain overly dependent on imported oil, and who were increasingly aware of the finite nature of 'non-renewable'

---

[4] See www.gbn.com.

resources such as oil" [40]. Scenarios being a recognized tool at Shell in the meanwhile, the company positioned itself accordingly, being more cautious with its investments than others in the business. The scenarios that have prepared Shell's management for the changes to come, are largely credited for the companies' ascension from 14th to the second oil company worldwide during the 1970s and 1980s (e.g., [40, 155, 140]).

In retrospective, it seems that such a development towards scenario thinking was bound to happen in a multinational oil company. Not only because the success of Shell (as of some of its rivals, too) depends on its performance on four continents and throughout different cultures (the planners having thus to take into account a very complex and diversified world), but also because in the oil business, the investments are large and long-termed. As noted in [40], "in the oil business you have to make enormous investments twenty years in advance – you have to build refineries that are not on line for years. So you're really just taking enormous risks with blocks of capital, all the time," the ability of planning several decades ahead thus being crucial.

Societal Responsibility through Scenarios

Of large interest in the context of this thesis is obviously the societal impact of the scenarios developed by Shell. Already during the original Wack scenarios that have foreseen the oil crisis, Shell took an unusual step. It informed the governments of the major oil-consuming countries about what they saw coming [161], putting its societal responsibility over the business interests. While, due to the fast development of the crisis, their advice did not have much effect in the end, it opened the way for a deep societal implication of Shell.

This growing implication originates in the insight that, besides the already discussed two levels, there's a third level to future scenarios. Such "third-generation" scenarios do not represent only "stories about the way the world might turn out tomorrow, stories that can help us recognize and adapt to changing aspects of our present environment," as Schwartz describes the second-generation decision scenarios. While in these latter an actor one can only react by trying to adapt now to changes expected in the future, the third-generation scenarios recognize that the future which will ultimately occur can actively be influenced, and are thus explicitly built with the purpose of influencing this path to the future. They are built to answer Roßnagel's question of "which future shall we want," by first presenting different futures, and then highlighting which measures will lead to all such individual future alternatives.

Perhaps the most striking example of such third-generation scenarios occurred in South Africa during the last years of the apartheid regime. The story is told by Betty Sue Flowers, who served as external editor for the Shell scenarios a few years after this episode, in an interview with Robbie Davis-Floyd [40]. When most companies retreated from South Africa due to the international sanctions against the apartheid regime, Shell has chosen to stay. Not only to do business, but also to propose third-generation scenarios. The scenarios, which became known as the "Mont-Fleur-scenarios,"[5] have been published in 1992 with the goal of stimulating the debate on how to shape the next ten years for South Africa. The result consisted of four scenarios with suggestive bird names: "ostrich," "lame duck," "icarus," and "flight of the flamingos." In the ostrich scenario, the white government continues to hide its head in the sand, not wanting to perceive the inevitable changes, but cannot fly away either. In the end, it will be forced to lift its head from the sand. In the lame duck scenario, the two parties agree to a slow transition, due to different reasons. While realizing the imperative of changes, the ruling minority fears an irresponsible government if these happen to fast. The suppressed majority fears a return to the further system if they do not make such compromises. Such a long transition with an imposed coalition is likely to incapacitate government. Indecisive politics, economic insecurity with low investments, constant or growing poverty are the results. The lame duck can not fly and has an uncertain future. The icarus scenario is one of macro-economic populism by a democratically elected government, which tries to repair too quickly the mistakes of the past and distributes lots of money to the formerly oppressed majority. Inflation, economic instability, business collapse follow quickly after the first months of euphoria. Like Icarus, the economy takes off quickly, flies high, but falls even quicker. To illustrate the real and not obvious danger of this scenario, the authors give several examples of Latin American and Caribbean countries that have succumbed to the attractiveness of an Icarus transition. Most of them, after the inevitable economic decline, have ended in an dictatorship. Finally, in the flight of flamingoes, the society as well as the economy take off as the flamingoes do: slow, but together, and continuously rising. The democratically elected government does some well-thought

---

[5] See www.gbn.com/ArticleDisplayServlet.srv?aid=455

social investments, especially to reduce violence, but always keeping the macro-economic parameters in mind. Because of the stability of the transition, the trust of businesses is gained, and investments start growing, creating jobs and reducing poverty. This does not happen overnight, but it happens eventually.

Interesting about the story is the large influence the scenarios seem to have had on all levels of South African society. According to Flowers, in the workshops organized by Shell, they brought together at one table ministers of the apartheid regime together with leading members of the ANC (African National Congress). The scenarios – being free of morals, just showing with economic and psychologic arguments how the society can evolve into different directions – provided a basis for discussion amongst people with very different ideologies. But not only on political level the scenarios have been discussed. After a 14-pages excerpt has been published in the two largest daily South African newspapers, it became a societal debate. Flowers says: "I observed what a difference the South African scenarios made. I heard preachers in their sermons referring to these scenarios, and ladies in the boondocks calling in on radio talk shows saying, 'I'm afraid we're going in the direction of Ostrich scenario.' It was important to see how a language of story could appeal so much, and become a language that all levels of society could enter into for the sake of democratic discussion." And about her discussion with South Africans of different origins: "They tell me all four stories, and say that clearly 'Flight of the Flamingoes' was preferable, and that everybody in the country knew the stories, and that those stories made it possible for people to understand that this decision leads to 'Flight of the Flamingoes,' and this one leads to 'Ostrich,' and this one to 'Icarus' – so it became not about your group wins versus mine, it became, 'This works – this works!' " [40].

4.2.4 Future Scenarios for Technology Assessment

From it's early days after World War Two, future scenario analysis has thus followed an interesting path. The military experts who first developed the technique have soon realized its societal value and started to use it to research societal questions, especially societal implications of technology. After having been adapted and refined for the business world, scenarios have again ended by having a larger societal reach. Due to their peek into the future, and inevitable broad scope, it thus seems almost as if scenarios invite their authors to regard societal questions, although their initial aim might have been different.

Within the context of this thesis, we are interested in the value of scenarios for the specific field of technology assessment. Except the possibilities that scenarios open in any domain, as discussed earlier, for technology assessment, they offer certain specific advantages. These are presented in [134] by Roßnagel, who has used scenarios among others to assess the future implications of the civil use of nuclear energy for Western Germany in the 1970s. As we have summarized in [30], scenarios offer the following specific advantages for technology assessment:

- Several different technology trends may be extrapolated and clustered into a single scenario. A broader range of possible consequences can thus be analyzed, since some of the consequences only show up when regarding the complete bundle of technologies and would not be evident if separately looking at the individual technologies. This is a particularly interesting feature for the field of ubiquitous computing, which inherently is such a technology mix.

- Various scenario techniques have evolved over time and they serve various purposes. Technology-driven scenarios, for example, allow the above-mentioned extrapolation of technology trends. They start with emerging or foreseeable technologies (such as ever improving location technologies) and analyze the possible consequences of these technologies. Problem-oriented scenarios, on the other hand, start with a specific class of problems (such as the everyday difficulties of physically impaired) and try to find means to solve or alleviate these problems through the deployment of a combination of upcoming technologies [134].

- Furthermore, in a scenario the different stakeholders can be analyzed together with their specific interests, so that a picture of future conflicts of interests can be drawn. Pierre Wack has been the first to discover the value of this feature, investigating the interests of the individual oil producing countries, and arriving thereby to surprising conclusions. For ubiquitous computing, the picture is likely to be more complex, since there are many groups of interests, e.g., producers of ubiquitous computing technology, companies who would want to deploy them, end consumers, governmental organizations, NGOs, or privacy advocates.

- Scenarios further offer the means for various interdisciplinary investigations focusing on aspects other than the main investigated issue. For example, by consequently listing all needed data collections in a ubiquitous computing scenario, a subsequent privacy protection analysis can ascertain the "price" (in terms of lost privacy) of the scenario. If the price is regarded as too high, other implementations can be investigated, that may offer the same advantages at a more acceptable price. A sociologist may analyze the same scenario from an inter-human interaction point of view, and so on.

4.3 Related Work: Ubiquitous Computing Future Scenarios

It becomes more and more evident that ubiquitous computing technologies will have a long-lasting, profound impact on the society. Since scenarios are a good instrument for analyzing the possible consequences, there have been quite some attempts over the last years to assess these implications through the use of future ubiquitous computing scenarios.

ISTAG Scenarios for Ambient Intelligence in 2010

Along with its efforts to promote the technologies of ambient intelligence, the European Union has also tried to assess the implications of smart environments. Its Information Society Technologies Advisory Group (ISTAG) has thus been called upon to realize a study, which has been published in 2001: "Scenarios for Ambient Intelligence in 2010" [46]. The study consists of four scenarios, all with a rather narrow scope. The first ISTAG scenario, "Maria – Roadwarrior", revolves around Maria, a business woman, who undertakes a business trip abroad. On her way from the airport over the hotel to her appointment, she uses personalized communication devices, adaptive rooms, and intelligent traffic systems, which help her complete several tasks, like finding a way around the traffic jam or projecting the slides finished during the flight. While the ideas presented in this scenario are not particularly thrilling, there would have been one opportunity to discuss the deeper implications of such systems – the potential for proliferation of free-market solutions, but also the risk of an extended "two-class society," consisting of the ones who can afford to pay for time-saving information and the ones who cannot – due to the relatively high amount of money payed by Maria to get navigated around the jam. However, this aspect is only superficially touched upon in the study. The second scenario, “Dimitrios and the Digital Me”, introduces Dimitrios, who – wanting to be free of annoyance during his lunch brake, but not miss any important call – is represented towards the outside world by his intelligent digital assistant, called “D-Me”. The D-Me can accomplish a multitude of tasks: understanding natural language, talking back in natural language while imitating Dimitrio’s voice, and understanding from such “discussions” if a call is of outmost importance (in which case it would forward it to Dimitrios after all), and it can have such conversations in several languages. Above all, it succeeds to negotiate with Dimitrios’ wife a delay for their phone conversation. The third scenario, “Carmen – Traffic, Sustainability, and Commerce”, describes a typical day in the life of Carmen Doe. Some ubiquitous computing artifacts are presented, such as smart refrigerators aware of their contents, cars that reduce their speed due to high smog values, or micro-payments in taxis. There are some points that could have been used for further societal investigations – e.g., vehicles automatically reducing the speed due to high smog values and not being “loyal” to their owners – but such discussion is again largely set aside. This scenario, too, assumes quite some amount of artificial intelligence, e.g., when Carmen’s apartment issues an alert for tomorrow’s demo and suggests Carmen to work from home. The fourth and last scenario resides furthest in the future and describes a meeting of environmentalists, which is supported by a so-called “ambient of social learning”. The ambient performs a multitude of tasks: it continuously changes the dynamic day-schedule, depending on current events, it can lead conversations with present and remote participants and collect their opinions, and finally it is able to “synchronize the mental states of the participants” [46]. How ambient intelligence is expected to perform all these tasks, is not explained, neither do the authors consider the potential negative implication of such omniscient ambient intelligence. The “ambient of social learning” is presented as a thoroughly positive development.

Summarizing, the four scenarios of the ISTAG study present either relatively trivial comfort gains (“Maria – Roadwarrior”, “Carmen – Traffic, Sustainability, and Commerce”), or they consider “intelligent” devices that always know their owner’s intents, and are able to perfectly understand and interact in natural language (“Dimitrios and the Digital

Me", "Ambient of Social Learning"). Such devices have been the dream of Artificial Intelligence for some decades now, and don't seem to be much more realistic today. As Lueg [106] puts it, such devices would need to "approximate" the user's behavior and wishes. Furthermore, the corresponding two scenarios do not take into account any negative implications such a penetration of life with intelligent artifacts could have, presenting them as genuinely positive trends.

TA Swiss Precautionary Principle

In 2002, the Swiss Center for Technology Assessment ("Schweizer Zentrum für Technologiefolgen-Abschätzung" in German, TA-Swiss for short)[6] commissioned a study on the longer-term effects of pervasive computing. The study has been realized under the guidance of Prof. Lorenz Hilty from the Swiss Federal Laboratories for Materials Testing and Research (EMPA) in St. Gallen, and included researchers from the following institutions: Institute for Futures Studies and Technology Assessment (IZT), Berlin, the University of Applied Sciences Solothurn (FHSO), Olten, the Institute for Business Ethics at the University of St. Gallen (IWE-HSG), as well as the Foundation for Research on Information Technologies in Society (IT'IS), Zurich. The original study was published in 2003 in German [74], an English translation was published two years later [75].

The authors "focus on the risks for human health and the environment" [75]. They further postulate from the start that the pervasive computing developments should comply to the precautionary principle and the paradigm of sustainable development. According to the authors, "the precautionary principle is used for dealing with risks in situations where there is no acute danger. Its purpose is also to minimize risks that may become evident only in the long term and to maintain a margin for future developments" [75]. Thereby, "the scale of a risk is assessed on the basis of the primary criteria of extent of damage and probability of occurrence. In so far as these variables are quantifiable, the scale of the risk corresponds to the mathematical product of the two primary criteria."

As presented in chapter 4 of the study, the authors have chosen four areas of life to investigate, since "it is difficult in the context of a study to investigate the prospects for realization and the effects of a tech-

---

[6] See www.ta-swiss.ch.

nological vision that is avowedly intended to affect *all areas of life*. Ubiquity, the penetration of ICT into all areas of everyday life, is the central idea behind the pervasive computing vision." Thus, out of nine potential life areas (housing, shopping, transport, food, health, clothing, information & communication, work, leisure), the authors have chosen four to look into: housing, transport, health, and work. They also regard three technological developments they expect to play outstandingly important roles for pervasive computing: wearables, digital information and home entertainment media, and electronic labels (i.e., RFID).

The study follows a complex scenario pattern. For every one of these four life areas and three technological domains, it pictures three future scenarios which differ from each other in the amount of pervasive computing technology penetration, since, as Hilty explains in an article that summarizes the study, "how the technology and the market will develop are open questions. We cannot predict with sufficient accuracy how fast and to which extent the technology will be taken up and how it will be used. Many predictions in the past have turned out to be wrong (e.g., nobody adequately predicted the success of SMS). Because of this reason, it is necessary to create scenarios describing possible paths of development, and base one's conclusions on the scenarios" [77]. There is the 'cautious' scenario which assumes restrained growth only, the 'high-tech scenario' "based on the assumption that everything that is technically and economically feasible will be achieved," and the 'intermediate' scenario residing in between the two and being regarded as most likely by the authors. Cautious, intermediate, and high-tech scenarios – all look at both a 5-year as well as a 10-year frame. The study has to deal with one further inherent uncertainty: the harmful potential of some of the risks, especially health-related ones, are difficult to predict. As Hilty writes, "for instance, there is still uncertainty as to whether exposure to non-ionizing radiation has non-thermal biological effects that can harm human health" [77]. In conclusion, the authors have screened, as they explain, for "potential risks." Since the term "risk" means "potential damage," the authors are thus looking at "potential potential damages," which seems a logical mistake. However, the first "potentiality" refers to the not yet conclusive knowledge whether the occurrence of some circumstance (such as non-ionizing) radiation implies a risk or not, while the second "potentiality" refers to the probability that the circumstance itself will occur or not. From these considerations, the authors conclude that *quantitative* risk assessment is not possible in the case of pervasive computing, and the study must regard *qualitative* risks.

On methodological level, the authors have used a three-step approach. In the first step, they have chosen the above-mentioned four potentially influenced areas, and they have developed three scenarios for each of them. The steps two and three have been done in a Delphi-resembling manner, during two workshops with experts from various interests groups and perspectives. Step two (first workshop) consisted of screening for potential risks, and resulted in a list of 23 health, environmental, and social risks. In step three (second workshop), a "risk filter" has been applied, to focus the study on the most relevant risks. Hilty shows in [77] the 23 initial risks, which then have been reduced by clustering to a number of 17, explains the risk filter, and shows how the resulting most relevant five risk clusters have been built.

Building the first systematic risk filter for ubiquitous and pervasive computing is probably one of the foremost contributions of the study, although it came more as a side-effect, as a tool that the authors needed in order to prioritize the risks. Since the risk filter is inherently subject to ethical values, the authors have used the concept of sustainability as defined by the United Nations' WCED (World Commission on Environment and Development), as well as criteria for risk filters from other sciences, such as chemistry. As a result, the study considers five criteria for prioritizing the risks [77]:

- *Socioeconomic irreversibility* – will it be practically impossible to restore the status from before the technology together with the risk have occurred? As an example, Hilty discusses car accidents that are a consequence of automobile technology, whose diffusion is irreversible.

- *Delay effect* – how long is the time span between deployment of the technology and the appearance of the damage? Obviously, "a long delay between cause and effect extends the time span of uncertainty and potentially leads to a situation with much higher damage because the cause has been given more time to spread before the effects were observed and countermeasures could be taken" [77].

- *Voluntariness* – is exposure to the risk voluntary? If not, the potential for societal conflicts is larger.

- *Fairness* – are the people benefitting from the technology the same who suffer from the risks? As in the case of voluntariness, if not, the potential for societal conflicts is larger. The current discussion about cars with high emissions might be such an example, where only their owners take advantage of the comfort and security they offer, but everyone suffers from the emissions.

- *Burden on posterity* – does the technology and its consequences affect the possibilities of future generations to meet their needs?

When applying the risk filter in the second workshop, the experts rated all previously defined risks with "low," "medium," or "high" in these five criteria. Only the risks that scored with "high" in at least two columns were further considered for the study.

As a result, the TA Swiss study raises numerous interesting issues that have then been further refined, such as the ecological impact of pervasive computing. It sets a frame of reference for analyzing environmental effects by identifying three kinds of environmental effects:

> "Direct (primary) effects on the environment will result from material and energy consumption in the production and use phases, including pollution caused by disposal of the resulting waste [. . .] Greater quantities and shorter service lives of components will most probably counterbalance or even outweigh the benefits from progressing miniaturization. The energy demand of the network infrastructure needed for Pervasive Computing might be as large as several percent of total power consumption if there are no incentives for using the technical energy saving potential.
>
> These primary environmental impacts of Pervasive Computing are to be seen in opposition to the secondary effects it provides in optimizing material and energy intensive processes, or in substituting pure signal processing for such processes (dematerialization). The potential environmental benefits from such secondary effects are considerable and can even outweigh the primary effects if, for instance, the increasing independence of activities from defined locations reduces traffic. But using these potential environmental benefits requires sufficient incentives to manage natural resources more economically. Otherwise, the growth in demand (tertiary effects) will

> counterbalance these savings. The experience gained thus far with ICT effects has shown that such a rebound effect occurs in most cases."

The rebound effect is explained by Hilty as "if a good gets cheaper in terms of its price or any effort necessary to obtain it, the demand for this good usually increases" [76]. An example for such rebound effects may be a smart navigation system, which calculates the quickest route according to the momentarily traffic situation. Compared to the static navigation systems used today, which do not take into account traffic jams, such a system would make journeys on average shorter, saving time, gas (because less is wasted in traffic jams), thus diminishing greenhouse gas emissions. In the terms defined by the authors of the TA Swiss study, these would be positive environmental secondary effects. However, the study argues that more often than not such positive secondary effects are outweighed by tertiary rebound effects: For example, if people know that journeys became faster, they will tend to plan longer journeys (using the saved time not for other activities, but for travel itself), which implies even larger amounts of gas consumption and polluting emissions. Another example has been researched by Hilty and his team. As they report in [76] after experimenting with different generation hardware platforms on which office workers had to accomplish similar tasks, the exponential progresses in computing hardware efficiency do not result in comparable work efficiency progresses for the people. In some of the measured cases there are some slight efficiency increases, in others the authors have even noticed a decrease in work efficiency.

The study has become the de-facto frame for researching the health-related and environmental effects of pervasive computing. It has also proposed the first systematic risk filter for ubiquitous computing technologies, which due to the sheer amount of possible risks within ubiquitous computing, is likely to become a highly relevant contribution for future research in the area. However, the study also has some weaknesses. For example, the authors state that the four life areas that have been chosen (housing, transport, health, and work) are the most likely to be influenced by pervasive computing. However, it seems that at least "smart housing" is a domain that will not have such a deep impact as other, not considered domains, such as shopping and the related dynamic pricing, which could have quite deep economic and therefore societal consequences. Furthermore, the chosen technological developments (wearables, home entertainment media, and RFID systems) are also questionable. While RFID systems are certainly one of the more prominent and controversial pervasive computing technologies, there exist more interesting technologies with larger impact than home entertainment systems and the still exotic wearable technology. Such example could be sensor networks, localization technologies, or new materials such as light-emitting polymers and smart ink.

Another point consists of the technologies envisioned in the scenarios. Since the study has been first published in 2003, the five-year time horizon of the scenarios can now be compared to reality. Almost all predictions of the intermediate scenario (expected to be closest to reality) have overrated the pervasive computing penetration. The study expects the computer in its present form to lose its dominance, being replaced by new types of mobile devices, or household devices routinely communicating over the Internet. Further, “in motor vehicles, from around 2006 the virtual safety belt becomes standard equipment. The car can automatically adjust its speed to traffic conditions. Excessively short safety distances and danger situations in poor visibility are communicated to the driver by speech output.” Even the cautious scenario overestimates in some instances the degree of pervasive computing penetration, for example by postulating that “from 2005 there are more Internet access points using multifunctional mobile terminals,” “refrigerators with a simple touch screen become popular. Their advantage is that they allow direct online orders from the kitchen,” or “in medicine, Pervasive Computing means among other things that medical practitioners in hospitals are networked” [75]. Why did the study overrate the amount of penetration of pervasive computing technologies? An answer could be the missing interdisciplinarity of the project in its starting phase. In the second and third phase of the project experts from various disciplines were present, among them researchers and developers of the technology itself. However, in the first phase, when the areas and technologies to be regarded have been defined, no pervasive computing developers were present.

Nevertheless, the TA Swiss study has been the first to propose a systematic approach for the thorough analysis of possible implications of ubiquitous and pervasive computing technologies, has been the cornerstone of numerous political and societal discourses, such as the “Pervasive Computing Dialogues” initiated by the Swiss “Risiko-Dialog” foundation,[7] and is still, five years later, the main reference for consequences of pervasive computing on health and environment.

The SWAMI "Dark Scenarios"

A more recent (January 2006) collection of ubiquitous computing scenarios has been developed by the same Information Society Technologies Advisory Group (ISTAG) of the European Union that had already developed the "Scenarios for Ambient Intelligence in 2010" [46] five years earlier. The new scenarios are part of the project "Safeguards in a World of Ambient Intelligence" (SWAMI), that has been promoted by the EU under its Sixth Framework Programme (FP6).

The emphasis of the new report, called "Dark scenarios in ambient intelligence: Highlighting risks and vulnerabilities" [5], is quite opposite to the early ISTAG scenarios. As the title suggests, the scenarios specifically highlight the potential drawbacks of the new technologies, while the previous scenario collection adopted a very optimistic (and, as argued above, often unrealistic) view of ambient intelligence technologies. Now, as the authors put it, "the SWAMI scenarios present visions of the future that we do NOT want to become realities. SWAMI has labeled them dark scenarios. The SWAMI scenarios [. . .] depict an undesired but realistic future that could emerge from the application of new ambient intelligence technologies."[8]

These "dark scenarios" are trend scenarios [109], extrapolating technological developments and showing possible futures. The time-frame is not prominently stated in the document, however, at two places dates in the future are mentioned: "Theft of data, 29 June 2015"[9], a recent event in the scenario and "As a former software engineer, Martin Schmitt (aged 77 in 2002), is familiar with technology."[10] Is is thus reasonable to assume that all scenarios have a 10-20 years time-horizon.

The scenarios have been conceived as to illustrate several types of possible negative consequences of an extensive ubiquitous computing deployment. The issues presented have been selected in a structured, multidisciplinary process, involving external experts at two points during the development. The participants rated potential impact and degree of probability for the individual negative consequences. By multi-

---

[7] See www.risiko-dialog.ch/Themen/Kommunikationstechnologien/263

[8] See [5], pp 13f.

[9] See [5], pp 48.

[10] See [5], pp 33.

plying the average scores of impact and certainty, the effects have been prioritized.

As a result, the problems highlighted from different viewpoints throughout the document are mainly privacy issues, the feeling of loss of control, the ever increasing dependency of ambient intelligence and the thus severe drawbacks in case the systems fail, the loss of digital identity, and the new types of criminality that could emerge. These problems are presented alongside two axes: One distinction is made between private (e.g., home) and public settings, the other between the problems affecting at one time individuals versus concerning a larger group at once. Privacy issues, for example, are shown on both a personal level in a private environment (the son circumventing faulty biometric access control mechanisms and discovering that his father likes erotic poetry and has bought quite expensive lingerie: “well, I hope it’s for mum”); a personal level in a public environment (a character being told by his friend-locator service that a good friend is in the neighborhood and goes searching for him, only to find him with an unknown younger woman and not with his wife); and on a societal level (information from a data mining corporation gets stolen and sensitive data from millions is on sale on the market).

Another development that is convincingly presented as being worrisome, is the increased profiling used by both law enforcement to narrow the group of suspects in an investigation, and private companies to spot potentially unreliable customers. In theory, it sounds genuinely positive that due to new technology it is easier to quickly narrow the suspects to a small group. The problem is that all persons in that group (but for the guilty one), are strongly suspected for a crime they did not commit. And such suspicions, even if later proved as being unsubstantiated, can easily bring disadvantages to the involved persons, as shown in one of the scenarios, where the character Paul misses a promotion in his security firm because the police has (on a very far-fetched suspicion) requested access to his data. Even afterwards, such persons will not know if there are still some records indicating that they’ve been suspected of a crime at some point and they will often have to fight with the blind windmills of bureaucracy to “clear” their name. The authors of the SWAMI scenarios hint here at examples such as the one of the 62-year old Dominican nun Sister McPhee, who has been harassed by security for nine months every time she entered the US, seemingly because a young Afghani man was using the name McPhee as an alias [142]. Such problems could indeed worsen as more and more computer systems gather and exchange data for profiling people. Names could suddenly pop up on some black list or trigger random alarms, almost without the possibility of reconstructing why the person aroused suspicion. Another scenario presents such a situation when the character Michael is not allowed to board a travel bus: “Apparently, some kind of data mismatch between his personal ID, the e-ticket and the information stored on the central server had caused the problem.”

Besides raising numerous relevant questions and illustrating them in an easily-accessible, pleasant to read manner, the SWAMI scenario collection has another major benefit. At the end of the document, for each scene depicted in the scenarios there is a thorough review of legal regulations in the EU and North America that would apply to the described situation. Thus, the legal context for the depicted negative implications can be quickly grasped and the reader can see whether the nature of the problems is so new that there is no regulation in place or whether the existing regulation is obsolete and can not apply any more to a world of ambient intelligence. As an example from the report, because of the continuing spreading of teleworking, traditional borders between home and work environment become increasingly blurry. However, this fact has implications on other fields, for example privacy protection. In today’s legislation, a person enjoys a great deal more privacy protection at home than in the office environment. Existing legislation, however, becomes increasingly outdated as it cannot cope with the employers need for more control over the teleworker and the need for privacy of, for example, other family members in the same household, who do not want to be spied upon in their home.

There are, however, also drawbacks to the SWAMI scenarios. First of all, the very fact that they are “dark.” The intention of the authors is easily comprehensible: to highlight only negative implications, so that legislators, system developers, and others can keep them in mind when plotting the future course for ubiquitous computing development. It represents nevertheless a disadvantage to have only one part of the story in the study.

Technology often is ambivalent, and ubiquitous computing constitutes by no means an exception. Aside from the comfort and efficiency gains they will probably induce, ubiquitous computing applications could also have significant positive societal implications, for example by allowing physically handicapped persons to lead a more independent life, or by enabling a higher degree of fairness throughout society, or by inducing major positive environmental effects.[11] In the SWAMI scenarios though, apart from comfort applications, the only mentioned positive societal effect is implicitly present since (elderly) people carry personal medical devices that are able to alarm the paramedics if the wearer has a critical medical situation. By not explicitly including a comprehensive analysis of positive effects into the study, the reader (politicians, interested public) can easily gain a biased image of ubiquitous computing. It appears as if the numerous listed societal problems are balanced only by some – on a societal scale rather unimportant – business interests.

Another problem of the report is that, although it has come a long way since the early ISTAG scenarios, it is still very technology-oriented. The trouble are not the futuristic gimmicks like PWC (personal wrist communicator), SAS (shopping assistant software), or HMD (health monitoring device), although they make way for involuntarily funny sentences like: "And thanks to the travel-assistance procedure of the AmI environment in our home, this time we even thought of recharging our PWCs and HMDs early enough." The problem are rather the still abundantly present AI-like intelligent systems, which know at all times what the user wants and means. Systems such as Maruja's SAS (shopping assistant software), for example, which is told to keep searching "until you find something really funny." A machine understanding the concept of funny so well that it can make comparisons and know what "really funny" is, seems quite improbable. Or such as Alvin, "the holographic embodiment of the company's embedded intelligence," which (or who?) immediately knows that the question: "Where's MacDonald?" refers to the company's vice-president, and not to the nearest McDonald's fast food restaurant or one of the other employees named MacDonald. Another example of environmental intelligence are the automated meeting minutes that are being generated by a computer. The "art" of keeping meeting minutes, though, is not writing down everything that has been spoken, with all unfinished half-sentences, jokes unrelated to the meeting subject, and sometimes too much talking on a side-branch of the main topic (which, probably, a computer with a good speech recognition system will be able to provide), it is rather to distillate the information and keep the essence (which, undoubtedly, a computer cannot do).

---

[11] Several other examples will be given in the next section.

Furthermore, while the study illustrates numerous issues that would constitute severe societal challenges, some of the issues presented seem unrealistic and very unlikely to happen, so there's not much of a point trying to deal with them. For example, people in the future society are too often depicted as relying exclusively on technology: "Maruja had forgotten to authorize an update of the location-based software and found herself walking in an area of the city frequented by drug addicts and criminals," or "The city's intelligent traffic system went completely mad and the resulting traffic chaos was the worst we've seen in more than 15 years. Traffic lights kept on alternating every five seconds at random, for almost an hour." Why Maruja should not know the no-go-areas of the city she lives in, or why the city administration is not able to deploy traffic policemen (as it happens today if the traffic lights fail) is left untold. Elsewhere, after a bus accident with several injured seniors, the paramedics arriving "set up a list of people with more serious injuries and those with private health insurance." While people with more exclusive health insurances do enjoy better healthcare on a day-to-day basis, it is highly unlikely that in an emergency they would be preferred over persons with cheaper insurance but more serious injuries. This would contradict the very basic moral principles that doctors comply to, be there new technologies or not.

Finally, several of the scenes depicted in the scenarios seem questionable from a technological or economic point of view. As noticed by Davies and Gellersen [39], one of the main problems with assessing the feasibility of thought-of future applications is that economic viability (i.e., will any user pay for such a service? Will anyone see a business model and offer the service?), and/or technological challenges (which can only be fully understood when developing a working prototype encompassing all major features of the system) are often ignored. An economic analysis, similar to the detailed snapshot of existing regulations relating to the scenarios, would have certainly benefitted the study. As for the technological feasibility, while scenarios are a good starting point, there should ideally exist a feedback loop from a working prototype, which usually raises new opportunities and challenges, to the scenarios.

4.4 Summary of the Existing Approaches

That in order to understand the longer-term implications of technology in general, and ubiquitous computing in special, pure technology-driven prototyping (as presented above and depicted on the left side of Fig. 4.1) is relatively useless, has been argued by many authors (e.g., [134, 39]). As already mentioned, numerous ubiquitous computing applications are restricted to the technical environment best known to their developers, and are thus narrow in scope [30] and unsuited to evaluate the societal impact of future technologies.

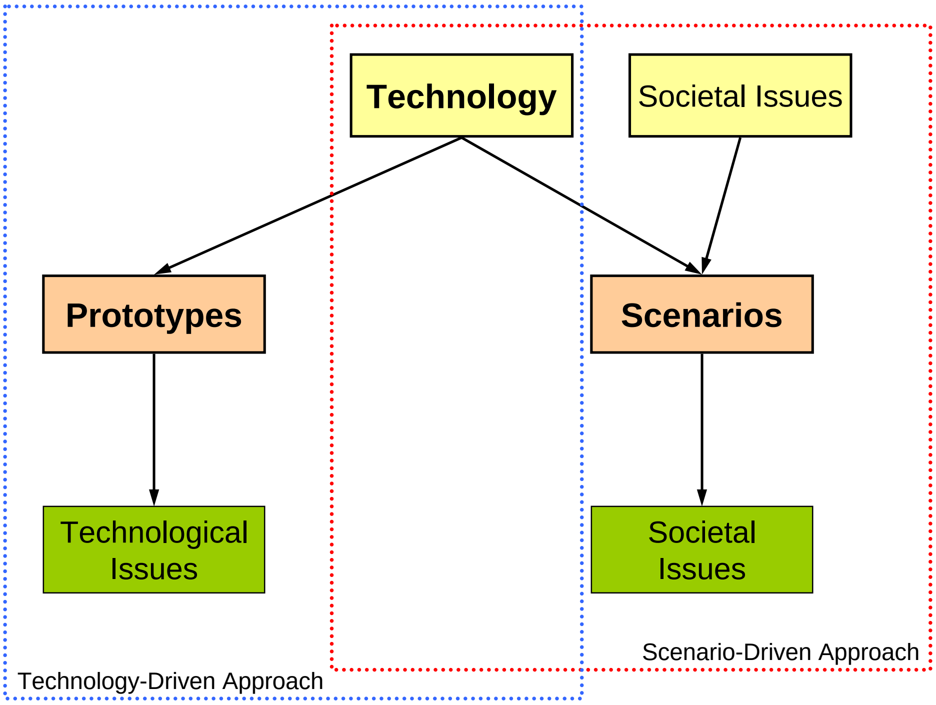

Figure 4.1: Two opposed ways to approach new technologies. On the left, prototyping cool gadgetry with minimal focus on societal implications. On the right, scenario-driven analysis of societal implications. The two are largely independent.

In contrast, the multidisciplinary scenario-analysis (as pictured on the right side of Fig. 4.1) offers several advantages: Due to the multidisciplinary approach, not only technological developments flow into the scenario, but also societal issues as analyzed by social scientists, such as societal needs, acceptance issues, economic feasibility, or existing legislation. Thus, such scenarios typically draw a more complex, differentiated, and realistic picture of what is to be expected from technological developments. After the scenarios have been completed, they also allow for new, cross-disciplinary analyzes, which may reveal new societal issues [30]. Thus, scenarios are both an instrument to present societal issues from different disciplines in a structured, condensed way, and a method that enables new societal issues to be discovered.

However, this second approach is also characterized by several disadvantages. Not only does it face the limitations presented earlier, which are inherent to the unpredictability of the future. It is also characterized by some more subtle shortcomings. While it does allow for the analysis of various societal effects, it often does so on a shallow, technologically naive level, which implies several limitations. Firstly, scenario-driven analysis does not provide the means to identify the technological challenges encompassed in the realization of the visions described in the scenario. This is more than a deficiency by itself. Without such technological "reality-check," there is a high risk for the visions and applications depicted in the scenarios to be partly unrealistic from a technological perspective, such as the ubiquitously present AI-like functionality described in the ISTAG scenarios [46], and still present (although much less predominantly) in the SWAMI scenarios [5]. Debating over the societal consequences of applications that will most certainly not exist is at best a waste of resources, often being downright counterproductive, due to the misleading of the public or political discourse. Secondly, as Davies and Gellersen observed, the more matured and realistic a prototype is, the more it allows conclusions for the economic feasibility of the vision behind it [39]. As they have noticed, many technologically feasible applications have not become reality due to the poor business plan behind them. However, only by developing a matured prototype can the technical challenges and thus the costs be realistically estimated. Otherwise, the same problem as before may arise: that scenarios deal with societal consequences of applications that will never exist, due not necessarily to their technical infeasibility, but to their missing economic viability. Thirdly, having only scenarios but no prototypes, inhibits also the back-link – from societal insights to technology development. If, for example, after the societal analysis of a scenario it seems reasonable to implement some of the application's functionality in a different way (to avoid some potentially damaging consequences, e.g., making the service more privacy-friendly), this is certainly an achievement. It represents, however, undoubtedly a much larger value to be able to prototypically test the new method and thus ascertain at which cost it comes, and which are the specific trade-offs.

5 Our Method: Scenario-Driven Prototyping

To overcome the limitations of the two presented approaches while keeping their specific advantages, we propose to use the paradigm of *scenario-driven prototyping*. This chapter starts by presenting the method in section 5.1. The method unifies the advantages of scenario analysis and technology-driven prototyping, but, more importantly, offers new advantages that only emerge from the synergy of the two. Section 5.2 then presents the historical background in which the method appeared and has been used – the Ladenburg Collegium, an interdisciplinary project researching the societal implications of ubiquitous computing technologies – as well as the five resulting scenarios from this project. The chapter ends by discussing in section 5.3 limitations and weaknesses of the novel paradigm, before the next two chapters will present the experiences and insights gained from two projects in which the method has been consequently applied.

5.1 The Novel Approach

As depicted in Fig. 5.1, the scenario-driven prototyping paradigm starts with the multidisciplinary development of scenarios. In this step, like in the scenario analysis, a multidisciplinary approach is followed, which complements the technological issues with know-how from social scientists. Thereby, relevant fields of application are identified with the specific opportunities they offer and risks they pose through the use of the new technologies.

However, aside from deriving societal issues directly from the scenarios, one important further step of the scenario-driven prototyping paradigm is to develop for each scenario one or more prototypes that highlight some or all aspects of the scenario. These are depicted on the left side of Fig. 5.1. By combining scenarios and prototyping, not only can the specific advantages of both approaches be exploited, there are

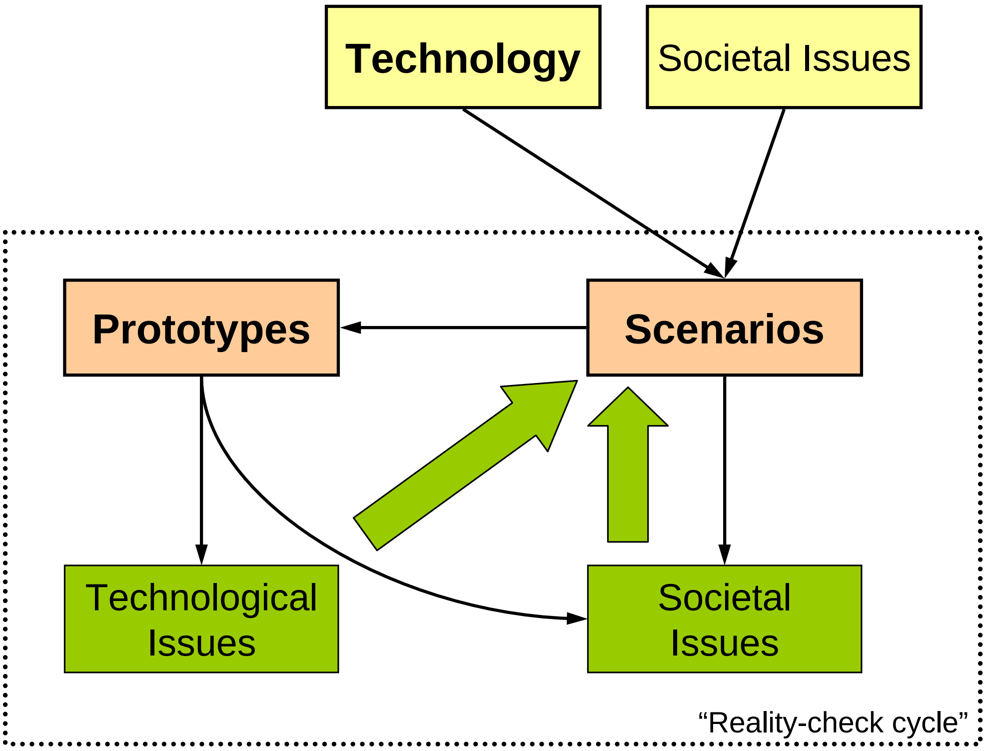

Figure 5.1: Scenario-driven prototyping, a paradigm that combines the specific advantages of both technology-based prototype development and scenario-driven implication analysis with new, specific benefits arising from the synergy of the two.

also new benefits that arise as results of this synergy. By having scenarios originating in a multidisciplinary dialogue, all advantages of the scenario-driven technology assessment analysis as presented in greater detail in section 4.2 are still present:

- The areas most likely to be affected in both positive and negative ways are better identified in a multidisciplinary dialogue. Thus, the scenarios (as well as later the prototypes) will be situated in relevant areas. To achieve this, both techniques of technology-driven and problem-oriented scenarios can be combined.

- Several technology trends can be clustered into a single scenario and extrapolated into the future to analyze their combined effects.

- The interests of stakeholders can be analyzed and probable future conflicts of interests shown.

- Moreover, the scenarios can be used to investigate further societal issues that have not been at the core of the original multidisciplinary dialogue. These novel societal issues are depicted in the lower right corner of Fig. 5.1.

On the other hand, the gains from having prototypes are also present:

- By developing prototypes, engineers can much better understand the technological challenges to realize real-world systems than if they would only theoretically try analyzing these challenges.
- As stated above, the more close to reality a prototype is, the better can the costs for such system be estimated.

But most importantly, there are some advantages that can only emerge by combining prototype development and scenario-based multidisciplinary analysis:

- By deriving the prototypes not directly from technological developments, but from the scenarios (which, in turn, resulted from a multidisciplinary dialogue), the prototypes are a great deal more likely to be situated in relevant areas and not be some simple gadgetry.
- Furthermore, some of the inherent limitations of technology only become apparent when trying to build systems. By detecting such limitations, the risk to have scenarios deal with unrealistic applications can largely be avoided.
- By being able to let potential users, stakeholders, but also social scientists "play" with a prototype, they are likely to realize new positive and negative issues of the technology, issues that they would not understand by only theoretically imagining the application. This is a great advantage of demonstrators, they make technology graspable, and people tend to more fully understand the implications of things when holding them in their hands. This direct link from prototype to novel implications is illustrated in Fig. 5.1 by the new arrow (as compared to the combined single approaches in Fig. 4.1) from "prototypes" to "societal issues".
- Finally, and arguably the most central synergetic benefit, is represented by the two thick arrows in Fig. 5.1. They stand for the conclusions that are drawn from either the technological feasibility or societal issues and flow back into the scenarios. With this

> back influence, a cycle results, where novel discovered societal issues can flow back into the scenario, altering it. Newly discovered societal opportunities can be included of the scenario, risks can be removed from it. By further propagating such changes to the prototypes depicting the scenario, conclusions can be drawn on whether there is a price to pay for including them. The price may be monetary, in terms of estimated production costs, but also in terms of different characteristics of the prototype, which may not be as strong in the novel scenario as they had been in the previous one. For example, by including a more privacy-friendly solution into a prototype, the system's security may weaken, or a higher level of trust is needed, which would in turn narrow the potential user basis. Such technological insights then flow back into the scenario, and the "price" that has to be paid can be judged from a social perspective. Then, again, these insights flow back into the scenario, and so on, the degree of refinement is theoretically limitless. Such a cycle of mutual insemination between technological and societal issues is only possible by including in the technology assessment procedure both scenarios and prototypes, as in our scenario-driven prototyping.

Summarizing, scenario-driven prototyping seems to be a more promising way for realistically assessing both societal opportunities and risks of ubiquitous computing technologies and thus strive towards a socially beneficial technology development [30].

5.2 Scenario-Driven Prototyping in the Ladenburg Collegium

Scenario-driven prototyping is not a mere theoretical construct, but the distillation of the practical experience from a three-year interdisciplinary project that tried to realistically assess possible long-term societal implications of ubiquitous computing. The project, which will be introduced in this section, brought together engineers and social scientists – i.e., representatives of both "fractions," the technology-driven prototyping and scenario-oriented analysis fraction – and had as outcome both ubiquitous computing future scenarios as well as prototypes underlying these scenarios.

While during the project, the process of scenario-driven prototyping as described in the previous section has not been theoretically conceptualized, the parts and bits composing it have been practically applied throughout the project's three years. The unique combination of practitioners and social scientists that have been unified in the project, and the intense cross-disciplinary discussions and fertilization that took place enabled the concept of scenario-driven prototyping, presented in this thesis, to emerge.

Section 5.2.1 introduces the project, its goals, as well as the project members and their responsibilities. Section 5.2.2 then presents the methods used within the project, which will explain the emergence of the concept of scenario-driven prototyping. Finally, section 5.2.3 presents the scenarios that emerged from the project.

5.2.1 Ladenburg Collegium: The History

The aforementioned interdisciplinary project has been the collegium "Living in a Smart Environment – Implications of Ubiquitous Computing"[1] for technological and societal assessment of ubiquitous computing. The collegium has been funded in its entirety by the Gottlieb Daimler- and Karl Benz-foundation[2] from Ladenburg, Germany – the foundation's goal being to fund basic research about the "interrelationship between Humanity, Technology and the Environment."[3] The collegium lasted for three years, between February 2002 and January 2005, and it unified researchers from different areas: computer scientists, engineers, lawyers specialized in privacy, psychologists, sociologists, economists, and philosophers.

By putting their various perspectives together, the participants tried to provide answers to questions such as those raised by Davies and Gellersen or brought up by concerned parts of the public: What are the areas that are likely to be affected by the large-scale deployment of ubiquitous computing technologies – in positive as well as negative ways? And what could be the long-term consequences? Out of this analysis, what conclusions for a socially sustainable technology can be drawn – so that positive developments are encouraged and those with large negative impacts are avoided?

The participants were divided into a so-called "inner circle" (people continuously working inside the project) and an "outer circle" (experts

---

[1] See www.smart-environment.de.

[2] See www.daimler-benz-stiftung.de/home/en/start.html.

[3] As stated at www.daimler-benz-stiftung.de/home/foundation/program/en/center.html.

Institution	Participants
ETH Zurich	Prof. Dr. Friedemann Mattern Vlad Coroamă
University of Kassel	Prof. Dr. Alexander Roßnagel Jürgen Müller Rotraud Gitter
University of Freiburg	Prof. Dr. Günter Müller Michael Kreutzer Adolf Hohl Moritz Strasser
University of Stuttgart	Prof. Dr. Kurt Rothermel Dr. Christian Becker Jörg Hähner
Fraunhofer Society	Dr. Dr. Norbert Streitz Carsten Magerkurth
University of Karlsruhe	Dr. Michael Beigl Tobias Zimmer
University of Rostock	Prof. Dr. Dirk Timmermann Matthias Handy

Table 5.1: The "inner circle" participants of the Ladenburg technology assessment collegium "Living in a Smart Environment – Implications of Ubiquitous Computing."

in various domains who have been regularly presented with the work progress and who provided regular feedback). The inner group consisted of participants from seven European research institutions, which are presented in table 5.1.

The collegium has been initiated and led by Prof. Dr. Friedemann Mattern from the ETH Zurich. Over the project's three years, the members of the inner circle have met ten times for one-day long debates in Ladenburg at the foundation's domicile. Almost all meetings also included one or more experts from the project's outer group, who provided scenario and prototype feedback from their particular research perspective. A typical Ladenburg meeting would thus comprise the presentation of the scenario progress, a practical demonstration of some new prototype that had emerged since the last meeting, followed by a discussion about all scenarios and prototypes from the perspective of the invited outer group members. The results of these discussions would then be considered for the next version of the scenarios, as well as for the prototypes. Table 5.2 presents the individual meetings in Ladenburg, their topics, and the invited experts for those topics.

Aside of these ten meetings of all collegium participants (comprising

Date	Meeting Topic	Invited Experts	Institution
19.01.2002	Project kick-off		
17.05.2002	Choice of relevant areas for the scenarios		
21.08.2002	First scenario presentation		
14.10.2002	Presenting future scenarios to the public	Dr. Michael Friedewald	Fraunhofer Society
		Michael Mikolajczak	Heinz Nixdorf MuseumsForum
24.01.2003	Loss of control	Prof. Dr. Christoph Hubig	University of Stuttgart
13.06.2003	Other future scenarios	Harald Vogt	ETH Zurich
		Patrick Keil	Technical University Munich
18.11.2003	Health and environment	Prof. Dr. Lorenz Hilty	EMPA St. Gallen
12.03.2004	Economy	Prof. Dr. Peter Welzel	University of Augsburg
25.06.2004	Privacy	Prof. Dr. Hansjürgen Garstka	Berlin privacy commissioner
		Dr. Marc Langheinrich	ETH Zurich
		Philip Robinson	University of Karlsruhe
16.12.2004	Project wrap-up; outlook	Dr. Werner Weber	Infineon

Table 5.2: The meetings of the Ladenburg collegium, their topics, and the invited experts.

both group leaders and the PhD students continuously working in the project), over the project's duration there have been more than a dozen meetings of the students only. These one to two-day workshops have been used to comment and critically review each others work, and to harmonize the efforts towards scenarios and prototypes. On PhD participants' level, it has been our responsibility to coordinate the scientific effort towards the completion of scenarios and prototypes.

The results of the collegium have been disseminated through different means. Aside of the numerous scientific publications that emerged from the project,[4] the collegium has also organized two public events to announce its findings. In March 2003, in Berlin, the conference "Total vernetzt – Szenarien einer informatisierten Welt"[5] aimed at disseminating the collegium's results about a future with ubiquitous computing technology (including the scenarios and prototypes) towards a more general audience, including politicians and other decision makers. The conference attracted an audience of over 150. The other event has been the two-day symposium "Der Computer im 21 Jahrhundert – Die Informatisierung des Alltags: Perspektiven, Technologien, Auswirkungen,"[6] organized in March 2005, right after the collegium's ending, in Zurich. This event attracted even more participants – some 250 – from research, industry, politics, and journalism. Finally, there have been other means of disseminating the collegium's findings as well. A short documentary by the Swiss National Broadcast Company SF TV[7] highlighted one of the project's prototypes (the Smart Tachograph, presented in chapter 6) and discussed the implications of such prototypes becoming everyday reality.

5.2.2 Ladenburg: The Working Methods and Experiences

What are the results of the Ladenburg collegium and which has been the working method used within the project? Why did this method enable the scenario-driven prototyping paradigm to emerge? These questions will be answered in the current section.

As it has been already briefly stated, the concept of scenario-driven prototyping has not been conceptualized during the Ladenburg process.

---

[4] Which are listed at www.smart-environment.de/publikationen.html.

[5] See www.smart-environment.de/berlin.html.

[6] See www.comp21.inf.ethz.ch.

[7] See on www.sf.tv/sf1/sfspezial/show.php?docid=20041201 the video entitled "Die kleinen Wunderchips."

However, the project synergetically brought together computer scientists, electrical engineers, and social scientists with the aim of answering the question about societal implications of ubiquitous computing as precise and from as many facets as possible. To accomplish this, the Ladenburg collegium has followed a process that can be roughly described by the following four steps:

1. In an interdisciplinary process, identify the most likely areas to be affected by ubiquitous computing, then develop future scenarios for those areas.
2. Implement prototypes and technology demonstrators to emphasize either entire scenarios or selected parts of them.
3. "Vertically" analyze scenarios plus prototypes from specific scientific and stakeholder viewpoints – with the goal of identifying novel societal opportunities and dangers, as well as technological challenges.
4. Let the insights from step three flow back into the scenarios, modify them accordingly, and analyze again at which costs the individual changes come.

Steps two to four can be repeated any number of times. In theory, every further such cycle should add more refinement and realism to both scenarios and prototypes, through the new issues that can potentially be discovered in each step. In practice, the value of the cycles would most likely sharply decrease over time, so that a small number of cycles should suffice.

The first step – elaborating the future scenarios – consisted of several iterations itself. At first, in an interdisciplinary dialogue process, computer scientists and engineers illustrated the possibilities of the ubiquitous computing technologies and gave a broad overview of the state-of-the-art research. Brainstorming with all the other scientists participating in the process, a number of 14 areas worthwhile for further investigation have been identified. Then, in a distillation process, five scenarios have been built out of these areas, by merging related areas into one scenario, by identifying similar questions, and by removing duplicates. Finally, input from stakeholders has been gathered for each scenario and then the five scenarios have been written down. The five resulting scenarios will be introduced shortly.

During the second – implementation – phase, as much input as possible that had been gathered in the first step from stakeholders and human sciences has been considered for the different prototypes. By trying to comply to the wishes of stakeholders and to avoid possible dangers identified by social scientists, several new technological issues became clear. In addition, the specific trade-offs between different possible designs (with diverse societal consequences) were vividly illustrated in this step.

Part of the third step – vertical scenario and prototype analysis – has been accomplished during the ten mentioned meetings of the project's inner circle members with experts from the specific highlighted domains. Other feedback has been gathered by presenting papers at scientific conferences or in journals. This step also involved a more thorough dialogue with stakeholders than in the first step. Being able to "play" with a prototype, the technology becomes more palpable for both stakeholders and scientists from non-technical backgrounds, so that they tend to realize some of the technology's potential better than from a theoretical presentation as in the first step. Thus, several new risks and opportunities have been identified in this last step. Also, after having already developed a prototype, the computer scientist and engineers are better able to appreciate the effort to modify the system so that it sustains desired attributes and it discourages debatable consequences.

Finally, all these different insights found, in a fourth step, their way back into the scenarios. As a result, the scenarios have been enhanced and large parts of them modified.

Furthermore, after scenarios and demonstrators were ready, we planned to have at least one more iteration of the aforementioned cycle, i.e., there should be a second multidisciplinary, stakeholder-involving analysis, centered around the prototypes. First, the insights and challenges from the development phase would be taken back to other scientists, who would then be able to analyze the scenarios again under the changed prerequisites. At the same time, if applicable, stakeholders should be involved in user studies of the prototypes. The feedback from both social scientists and users should then influence back the scenarios and subsequently the demonstrators, if necessary.

As Fig. 5.2 shows, the four steps presented above perfectly fit into the scenario-driven prototyping paradigm, which comes as no surprise since the method distillates precisely these practical experiences. However,

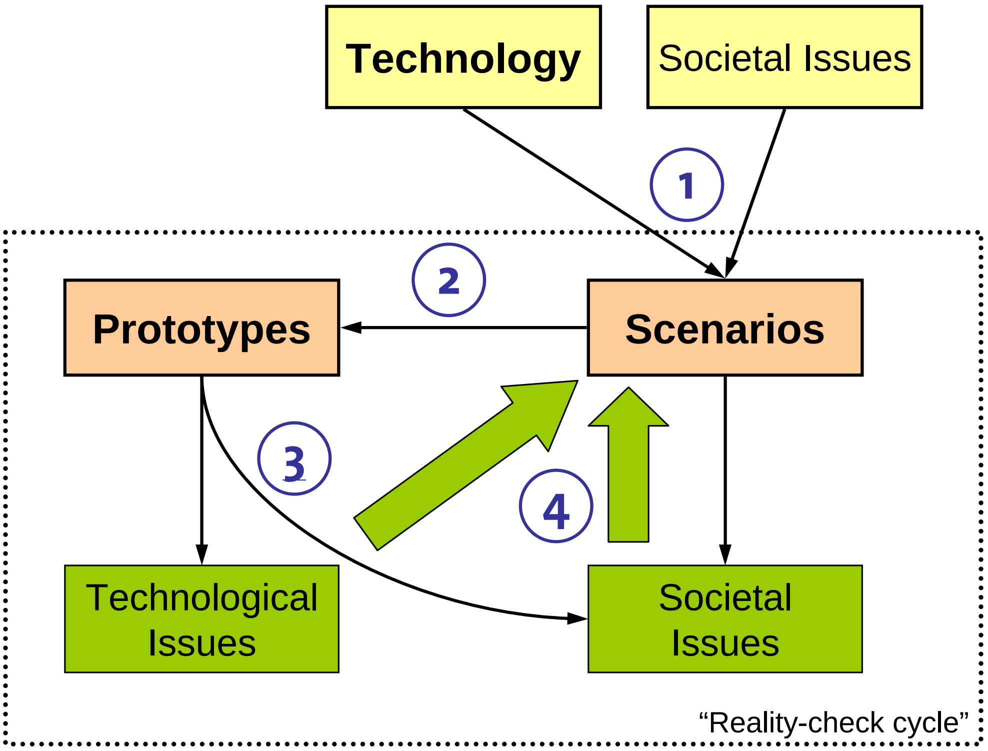

Figure 5.2: The four steps of the Ladenburg scenario-driven prototyping method. First, developing scenarios in a multidisciplinary process. Second, implementing (part of) the scenarios as prototypes. Third, understanding new technological and societal opportunities and risks. Fourth, letting these insights flow back into the scenarios and modify them accordingly. The steps two to four can be repeated several times.

these steps have not been so clearly separated, nor have they been precisely conceptualized during the project.

In practice, the theoretical scheme from Fig. 5.2 has not been dogmatically followed. The iterations have been more numerous, however of lesser amplitude. The project has lasted for three years, of which the scenarios took about the first year to be accomplished. Nevertheless, since the work for the first prototypes started after around six months into the project, the early results from the first "technology check" cycle have been included into the scenarios. Moreover, during the project's three years, at all of the ten plenum and over a dozen of focus meetings, the scenarios and the prototypes have been discussed and continuously improved. Since the results from these continuous refinements could not be reflected into the scenario collection any more, they are scattered among the numerous papers describing the individual prototypes and the lessons learned out of them.

Conceptualizing the process, proposing the paradigm of scenario-driven prototyping, and underlining the concept by applying it in two domains (presented in chapters 6 and 7, respectively) are the main contributions of this thesis, and will be presented as follows: The five scenarios resulting from the Ladenburg project process are briefly presented below. The two of them that have been further pursued by us (including the whole scenario-driven prototyping cycle: implementation as prototype, user studies, feedback gathering, modification of scenario and prototype) will be presented in chapters 6 and 7, respectively.

5.2.3 Ladenburg: The Scenarios

The scenarios developed within the Ladenburg project have been published in the report “Leben in einer smarten Umgebung: Ubiquitous-Computing-Szenarien und -Auswirkungen” [30]. Several areas have been depicted in these scenarios, in an attempt to cover different typical life situations, but also to show some new applications enabled through the use of ubiquitous computing technology. Thus, the scenario collection comprises three technology-driven and two problem-oriented scenarios. The shopping in a supermarket, a business meeting including the journey of external meeting attendants, and the analysis of public and private transportation, constitute the daily situations covered by the technology-driven scenarios. The logistics of medical equipment in a hospital and a vastly enhanced everyday support for the blind are presented in the two problem-oriented scenarios, which look at specific, narrow, but nowadays unsolvable problems, and how these could be tackled by the use of ubiquitous computing technology.

The scenarios have all been presented according to the same, fourfold principle: A first part describes some daily process such as the shopping in the supermarket from today’s perspective, highlighting the disadvantages that can arise (e.g., out-of-stock problems, forgotten goods in the storehouse, queues at check-out, etc) as well as the limited ways to counter these problems nowadays. The second part takes the same process into a ubiquitous-computing-pervaded future, and shows how many of these problems will be more elegantly solved. To not follow a technological naive path as done, for example, in the ISTAG collection of scenarios [46], the third part of each scenario comprises a detailed description of the technologies needed (both in terms of devices and infrastructure) for the scenario to be feasible. Thus, at any later point the degree of realism of the scenario can be reassessed. Finally, the scenarios have been analyzed from the specific perspective of the participating experts of the collegium's inner and outer circle, identifying the scenario's most relevant societal consequences. A preliminary analysis (since it does not yet include the more profound reflections from the prototype analysis) of positive and negative implications forms the fourth and final part of each scenario presentation.

The scenario collection finishes with a wrap-up of 14 of the most relevant implications. These are the ones recurrently appearing in several scenarios (i.e., expected to show up in different life situations). They include socio-technological acceptance aspects (loss of control, overly reliance on systems, usability and manageability), economic issues (economical interests involved herein, new economic paradigms), social compatibility questions (fairness, digital divide, universal access to information, information reliability), privacy questions, as well as more general problems with the deployment of ubiquitous computing systems (needed ontologies, providing user feedback).

The five scenarios will be briefly introduced here. The two in which we have applied the scenario-driven prototyping method to assess the specific implications of ubiquitous computing will be presented in larger detail over the next two chapters. They are the *shopping* and the *independent living for visually impaired* scenarios.

Scenario: Shopping

The first scenario shows the profound impact that ever cheaper and smaller sensing, computation, and communication devices (that could thus eventually be integrated even in the cheapest of supermarket items) could have on the process of shopping. It depicts a future supermarket in which the products have some basic "self-awareness," being aware of attributes such as their best-before-date, whether the conditions in which they have been transported and deposited correspond to their needs, weather, time, or how many other products there are around. Depending on all these attributes, every item continuously adjusts its price to fit the momentary demand and maximize sales. The concept of a generalized dynamic pricing is at the core of this scenario.

Scenario: Logistics in a Hospital

The area highlighted by the second scenario is the planning and distribution of resources in a hospital – these can be large, expensive and sparse instruments like x-ray machines, but also doctors and nurses. Nowadays, hospitals have to cope with numerous problems in the planning of their resources, for example the insecurity over the number of emergencies that will come in (which disturb the normal flow), or lack of knowledge of the whereabouts of personnel. Other, more serious problems occur today as well. One of the most prevalent is the lack of comprehensive and ubiquitously available knowledge about the medication prescribed to each patient, or the medication that has been taken by each patient. Due to the paper-based, error-prone distribution of such information, every now and then fatal errors occur. The scenario shows how ubiquitous computing could help to largely eliminate such problems and inefficiencies.

Scenario: Living and Working

The third scenario shows everyday problems arising from the lack of extensive, up-to-date information in unknown environments. A business traveler, who arrives for a business meeting in an unknown town, gets confused by a bus stop that has been moved due to roadworks. Then, trying to make use of the few free hours he has at his disposal, he misses a guided tour of the local museum by five minutes. He could have made the tour by having a coffee afterwards and not before, if he only had knew the hour the guided tour starts. The future part of the scenario shows how through the use of ubiquitous computing technology all this information becomes available at one's fingertips.

Scenario: Travel and Transport

The fourth scenario shows the economic, ecological, and time efficiency advantages that are likely to arise in both public and private transport in a future pervaded by ubiquitous computing technology. In the future part of this scenario, there is a high degree of networking and information exchange between numerous systems. The alarm clock is connected to the traffic information, so it can react flexible and ring earlier or later, depending on the traffic situation. The car's navigation system is connected to the traffic information as well, taking the current traffic flow into account when suggesting the route. All vehicles are further connected to the municipal park lots managing system, announcing when they free a parking lot, information that is at once broadcast to all other vehicles driving around town. Similarly smart solutions govern public transportation as well: Mobile assistants quickly guide the traveler through an unknown railway station. He does not need to queue for tickets any more, since he is registered when entering and exiting the train and the precise amount for the trip is deducted from his account. Several other types of information, such as train delays, location of free seats, or connecting trains, are also available on his mobile assistant.

Scenario: Independent Living

Finally, the last scenario shows the possible improvements in the lives of visually impaired that become possible through the new technologies. Nowadays, visual impairment comes with a distinctive characteristic: the blind person is in need of guidance and assistance. The daily shopping in the supermarket, for example, appears to a blind person as a maze with thousands of items, feeling all the same, impossible to determine by their shape, spread over hundreds of identical shelves. Visually impaired people will therefore only go shopping to their local supermarket and buy only few products in well known locations. The scenario describes a future mobile assistant for the blind and visually impaired, that reveals them their surroundings as they walk through the city, or in a supermarket. The mobile assistant using object recognition is combined with a speech interface, which enables the visually impaired to interact with the environment, thus providing them with a much larger degree of independence than today.

5.3 Limitations and Weaknesses of the Method

Scenario-driven prototyping offers, through the tight coupling of prototype development and scenario analysis, some unique advantages, as the next two chapters (which present two projects developed according to this method) also account for. Nevertheless, for several aspects of technological and societal assessment, the method is at the same time either constricting, or does not make any contribution where other methods do, or can only contribute when complemented with other already existing techniques. These limitations are presented in the current section.

5.3.1 Limited Time Horizon through Prototypes

A first, rather obvious drawback is the temporally limiting effect of the method, due to the prototypes it requires. Classical scenario analysis has no theoretical time limit for the peek into the future. As presented in the last chapter, the first Delphi study [64] looked no less than 50 years into the future. While the results of this half-century attempt to understand the future are rather questionable, as we have already argued, a long time horizon is sometimes necessary, as the original Shell scenarios [161] have proven. Some technology assessment scenarios, too, only have value when the time frame is long enough, as the example of civilian use of nuclear power shows. Even for ubiquitous computing, where the technologies are vast and rapidly changing, Mattern suggests that it is possible to identify the broad technological trends for the next 10–15 years [114], which may be a longer peek into the future than allowed by our method with its request for prototypes. Enforcing the coupling of scenarios and prototypes means that the scenario cannot be situated too far in the future either, since some of its relevant parts have to be prototypically implemented.

This drawback, however, has to be put into perspective. As Mattern [114] further argues, the technological trends of ubiquitous computing are almost impossible to predict for more than these 10–15 years in advance. Even Moore's Law [117], one of the main foundations of the ubiquitous computing visions, is only expected to continue for this period. Whether and how the miniaturization of microelectronics will go on beyond this horizon, when transistors will already have reached a size of only a few atoms, is yet unclear. Thus, unlike nuclear energy or other slower evolving technologies, attempting technological predictions for ubiquitous computing for more than these 10–15 years into the future would be a fruitless attempt, regardless of whether prototypes are to be implemented or not. Since typically some years pass between the first availability of a technology (when it can already be used for prototypes) and the moment it becomes ready for the market (when applications can use it and potential consequences start to develop), scenario-driven prototyping already covers a part of this period.

Furthermore, even if some technologies are not available yet, for a prototype they can often be substituted with other technologies providing similar services, or the same service in a more restricted environment. If, for example, a project wants to analyze a future application that can track its assets with a precision of a few centimeters, as promised by the European Galileo positioning system scheduled to operate in 2013,[8] it could build a prototype with less precision using the GPS positioning system. Or it could use the same precision of a few centimeters already offered by a UWB (ultra-wideband) [168] indoor positioning system, which, however, only works for relatively small indoor environments.

Finally, there is one more differentiation to be made: While the technologies depicted in the scenario cannot be too futuristic since they have to be prototypically available, this limitation does obviously not apply for the implication analysis. The consequences of technological deployment can be analyzed with scenario-driven prototyping, as with any other method, for a virtually unlimited period. Since, as Hilty [77] argues, risks of technology sometimes appear only after the technology has already been deployed for a long time, this look beyond the moment of deployment will often be necessary. Scenario-driven prototyping not only does not hinder this, but it might even encourage the analysis to confidently look further into the future, aware of the strong technological grounding in the present.

5.3.2 Long Time to Accomplish the Analysis

Scenario-driven prototyping is further likely to require more time to reach its results than other scenario methods. This stands to reason since the method involves a multidisciplinary dialogue, another dialogue with stakeholders when appropriate, the writing of scenarios, the development of prototypes, and usually more than one cycle of all of the above. The affirmation is also supported by our experience from the Ladenburg collegium, where the whole process, from the first dialogue to the final versions of scenario and prototype, has taken for the two projects two and a half and three years, respectively.

This seems a more serious drawback, especially due to the complex and dynamic nature of ubiquitous computing (attributes which are, at the same time, among the main motivations for our method). The question arises: Can a method, which may provide more precise answers than others, but only after a few years, be of value in such a dynamic field? While we cannot provide a definitive answer, the empirical ev-

---

[8] See www.esa.int/esaNA/galileo.html.

idence from the two projects that will be presented in the subsequent two chapters suggests that scenario-driven prototyping is a valuable tool despite the longer time needed to complete. In both projects we could – by following the method – provide valuable technological and societal contributions. Despite the time needed to complete the analyzes, our contributions have been made well ahead of the moment when a more general societal discussion started about the applications fields and – in one of the two cases – industry and politics started to show interest in the area.

The explanation for this seeming paradox (in a rapidly-evolving field, a method needing more time to complete than others can nevertheless contribute early enough to the field's understanding) might lie in the quest for relevant areas, scenarios, and prototypes that scenario-driven prototyping has been conceived for. While other methods, if employing a relevant future application, with strong business model and societal value, can provide good technology assessment results earlier than with our method, often these very application domains are not found. Our method, while slower, runs a smaller risk of analyzing irrelevant application domains.

5.3.3 No Contribution to Scenario Techniques

A further limitation of scenario-driven prototyping is that it does not contribute in any substantial way to the science of technology assessment through scenario analysis. Where, for example, the Delphi method introduces successive rounds of expert interviews together with median and quartiles result presentation [64], Pierre Wack differentiates between first and second-generation scenarios and introduces the empirical insight about the ideal number of scenarios [161], or Hilty introduces a risk filter for ubiquitous computing technology assessment [77], we do not provide any advances.

This, however, has never been our intention. Since this is a computer-science thesis from a ubiquitous computing professional, we only wanted to stress the importance of bringing together scenarios and prototypes, and of letting the insights gained from their analysis to flow back into a new version of scenarios, which is what the experience from the ubiquitous computing technology assessment projects where we applied the method has taught us. This tight coupling of scenarios and prototypes is not only important for the technological grounding of scenarios so that they do not become detached from technological realities, but also for the advancement of the ubiquitous computing field itself, since looking at more likely applications means discovering and approaching more relevant technological challenges.

Instead, our method lies orthogonally to classical technology assessment research and can be combined with any of the methods proposed there. For example, for the first step in our method (the development of the first version of the scenarios as depicted in Fig. 5.2), we only postulate that it should be accomplished in a multidisciplinary approach and, if applicable, that it should involve stakeholder interviews as well. This multidisciplinary analysis could, however, consist of a Delphi-style analysis, or of the development of second-generation scenarios in Wack's terminology. Our third step (identifying technological and societal issues from scenario and prototype, see also Fig. 5.2) could consist, for example, of steps two and three of Hilty's approach: identifying, in a multidisciplinary dialogue, risks and opportunities, and applying the risk filter, respectively.

5.3.4 Unclear Applicability Outside Ubiquitous Computing

Finally, and related to all points above, while our method has proven its use for ubiquitous computing, it is unclear whether it could successfully be applied outside this field. There are some areas, such as the civil use of nuclear energy, where prototypes cannot be built due to physical laws. In others, prototypes might have no value since they cannot catch any features of the large-scale technology. A lab "prototype" of an offshore oil drilling platform, for example, would probably not allow any conclusions as of the ecologic impact of the former. For yet other technology assessment areas, especially such analyzing one technology only, even if prototypes are feasible, they would probably not generate more profound insights than a standalone, thorough scenario analysis.

From today's technological areas, we do not see any other where scenario-driven prototyping would bring such an added value as for ubiquitous computing. No other field seems to be so broad and rapidly evolving at the same time as to need to be forced into the confines of a tight coupling between scenario and prototype development. Maybe it is precisely because of this uniqueness of ubiquitous computing that a novel method for its technology and societal implication assessment was needed.

5.4 Summary

In chapters 2 and 3, it has been argued that ubiquitous computing is being pushed both technologically and economically. It has further been shown that these novel technologies will most likely influence many aspects of our everyday lives. Hence, it is imperative to look as early as possible to the possible societal consequences. However, due to the sheer amount of upcoming technologies, possible applications, and areas that could be affected, it is a challenging task to pick the domains with expected high societal impact.

Thus, the topic of the current chapter has been to introduce a structured approach for a realistic assessment of relevant ubiquitous computing implications. It is built upon the related work discussion from the last chapter, which presented the two approaches predominantly used so far for technological assessment: the technology-driven prototyping, a path typically taken by engineers to be able to assess the technological challenges ahead of them, and the scenario analysis, typically used by social scientists to illustrate possible advantages and pitfalls of upcoming technologies. Although such scenario analysis presents several advantages, related ubiquitous computing future scenario collections also point to one of the largest drawbacks of this technique: not being anchored in a profound understanding of the technologies involved and sometimes ignoring the economic context as well, future scenarios easily run the risk of dealing with more science-fiction-like rather than realistic future technologies, unlikely applications, and thus irrelevant societal consequences.

The chapter has introduced our novel method, scenario-driven prototyping. It has been argued that this paradigm succeeds in combining the advantages of both classic methods. Moreover, it also offers some unique advantages, which can only arise through the blending of scenarios and prototypes, a blending that necessarily has to be accompanied by an intensive dialogue between engineers and social scientists. The ideas conceptualized within the scenario-driven prototyping method emerged as result of the experiences gathered within a three-year interdisciplinary project, the Ladenburg collegium “Living in a Smart Environment – Implications of Ubiquitous Computing.” The history of the collegium, its goals and participating scientists have been presented, together with a description of the multidisciplinary dialogue taking place throughout the project and its public results. Similar to the practice during the Ladenburg collegium, scenario-driven prototyping consists of four main steps: First, developing scenarios in an interdisciplinary process which takes into account both technological possibilities and limitations, as well as societal needs and boundaries. Second, prototypically implementing parts of or the entire functionality described in the scenarios. Third, in a novel multidisciplinary analysis of the prototypes – an analysis involving engineers, social scientists, and, where applicable, stakeholders – discover novel issues and opportunities of the future vision. Fourth, letting the insights from this step flow back into and modify the scenarios. Steps two, three, and four can thereby be repeated several times.

Finally, before the chapter ended with an analysis of limitations and weaknesses of the scenario-driven prototyping method, the working method within the Ladenburg collegium has been presented and the five scenarios of the collegium have been sketched. The two scenarios that have been further pursued by us to demonstrate the usefulness and significance of the scenario-driven prototyping paradigm will be thoroughly presented throughout the next two chapters.

These next two chapters represent a proof of concept for the paradigm of scenario-driven prototyping. We have applied the method to two application domains of ubiquitous computing technology, personalized prices and orientation aids for the visually impaired. The chapters will follow a path parallel to the steps of scenario-driven prototyping. For the sake of simplicity and readability, these chapters will not chronologically follow the whole iterative process as described above, with all its numerous small cycles. They will be built according to a simplified scenario-driven model with one cycle only, thus consisting of four major steps. The first parts will comprise a motivation for the importance of the respective area, which can be seen as a preparatory – and typically interdisciplinary as well – step for the scenarios.[9] The second parts will include a detailed description of the respective scenario as well as an analysis of those societal issues that have been acknowledged from the start (i.e., before the prototypes have been developed) as being relevant. The third parts will expose the implementation details of the respective prototypes. Finally, the fourth parts of each chapter will present the issues that have been discovered during the “vertical analy-

---

[9] The reader will remember that, above all, we do not want to waste resources to analyze societal implications of technology or applications which are so far-fetched that they will most likely never become reality. This very first step is thus of outmost importance and has been regarded as such during the Ladenburg process.

sis" of the respective scenario and prototype from the different scientific and interest group viewpoints.

6 The Smart Tachograph – Individual Accounting of Traffic Costs and its Implications

As argued in the previous chapters, one of the foremost consequences of ubiquitous computing technology will be the ability to model the physical reality in computer systems more closely than ever before. Cheaper, tinier, ever more numerous sensors will observe increasingly larger parts of the physical world, communicate their observations wirelessly, thus being able to share the gathered data among them and with the background computing infrastructure. At the same time, storing and processing the growing amounts of data will be no problem either, due to the dramatically decreased costs.

From the early stages of the Ladenburg project it became clear that such growth will have a deep impact on the economy, changing existing business processes and enabling entirely new models. This trend had already been recognized as early as 2001 as a natural extension of the nowadays "just-in-time" economy, slowly yet steadily transforming the "new economy" of the late 90's into what has been coined as "now economy" [141]. Section 3.1 has presented several of the new or improved business models that have been envisioned as consequences of the propagation of ubiquitous computing. Within the Ladenburg scenario collection, the shopping scenario[1] focuses on some of these aspects, most prominent among them being the proliferation of individual, dynamic pricing, which is context- and behavior-dependent. *Individual and dynamic pricing* refers to products (most of which we have grown accustomed to having fixed prices) that start exhibiting a much more dynamic price structure, their momentary price depending on numerous factors of the surrounding real-world.

This paradigm is presented in the Ladenburg shopping scenario by means of a supermarket, where all products continuously adjust their prices. While analyzing the scenario together with economists, we

---

[1] See [30], pages 11–31.

learned that this example, while underlining a relevant idea – the expansion of individual, dynamic pricing – is rather far-fetched in terms of economic feasibility and public acceptance. In subsequent discussions, the dynamic pricing model has thus been refined and presented within a new, to a large extent more realistic domain – highly personalized accounting of traffic costs. According to the principle of scenario-driven prototyping, in order to avoid the development of irrelevant prototypes, not a supermarket with variable product prices, but a system for the individual and dynamic accounting of traffic costs has prototypically been developed.

The chapter is structured as follows: section 6.1 presents the paradigm of dynamic pricing, along with its advantages and drawbacks, and the drivers behind it. Section 6.2 summarizes the shopping scenario together with its subsequent interdisciplinary analysis, which eventually led to the modified scenario. Section 6.3 shows the developed prototype. Section 6.4 presents the results of the vertical analysis of scenario and prototype by scientists and stakeholders. Finally, section 6.5 presents the related work.

6.1 Motivation: Price Discrimination, Expiration of Products

For the choice of the shopping scenario as one of high relevance for ubiquitous computing technology assessment, two economic factors have been decisive: the concept of price discrimination and the problem of depreciated goods in retail. They are presented in this section.

6.1.1 Price Discrimination

Price discrimination is the "practice of selling a commodity at different prices to different buyers, even though sales costs are the same in all of the transactions" [19]. The vendor's incentive when using price discrimination is to exploit as much as possible from each consumer's readiness to pay for a certain product (which, obviously, varies from customer to customer). The concept has first been defined by Pigou in 1920 [124], who has also defined today's widely accepted three types of price discrimination: In first-order price discrimination, also called "perfect price discrimination," "the price corresponds to each consumer's readiness to pay, provided that all variable costs are covered" [121].

It is relatively straightforward to see that this maximizes the vendor's profit: he doesn't make losses for any product (since at least the variable costs are covered), he doesn't lose any customer that is prepared to pay more than the product's variable costs, and he gets from each customer the maximum amount of money he or she is ready to pay.

However, the first-order price discrimination is merely a theoretical model and unrealistic in practice: firstly, since it would imply the seller's ability to know for each customer the maximum price he or she is prepared to pay [143]; secondly, because it would require "some means to discourage discount customers from becoming resellers and, by extension, competitors" [167]. Thus, more relevant for the practice are the second- and third-order price discriminations. In the relatively straightforward, well-known second-order price discrimination, the price varies with the quantity sold, larger quantities being available at a lower unit price. Third-order discrimination is the more subtle and challenging dividing of customers according to geographic location or customer segment [167], in order to exploit the average readiness to pay for every particular segment.

In practice, third-order price discrimination can often be observed, especially in monopolistic [157] and oligopolistic [152] markets. Empirical studies have, for example, proven third-order price discrimination in the European car market [158]. Another well-known example of price discrimination are prices for airline tickets, where advance-purchase discounts and Saturday-night stayover requirements are just some better-known among the numerous tools to charge different categories of customers (e.g., business travelers vs. tourists) individual prices for the same good [150]. However, not only for such relatively high-priced items price discrimination can be observed. For several rather cheap supermarket items, such as cereals, paper towels, coffee, fishes, or ketchup, price discrimination undertaken by the oligopolistic supermarket chains has been determined [152].

Obviously then, price discrimination is an important concept to look at, because of the strong incentives for vendors to pursue it. If a particular technology is expected to boost its usage, the technology will most likely be deployed. This, however, is not the only reason for the concept's importance. A larger degree of price discrimination would also imply overall welfare effects, and thus gain a deeper societal meaning. Indeed, as Varian [156] has shown, price discrimination increases overall social welfare due to the fact that it increases economic output, which has positive effects both for the vendor, but also for employment and salaries.

With the increase of computers and Internet usage, a tendency towards more price discrimination has already been observed. As Skiera argues, "…this ongoing digitalization, of both goods and processes, leads to a shift of weight in the relation of variable to fixed costs in favor of the latter" [143]. In other words, for digital goods and processes, the fixed production costs dominate more and more over the variable costs, which tend to decrease towards zero. For a large software product, for example, such as the Acrobat Reader or Microsoft Office, producing the first item implies a large amount of fixed costs, while the variable costs for producing more copies are almost negligible. For the distribution of goods over an Internet shopping portal, the lion's share is also comprised of fixed costs, even when the goods are not digital (e.g., software), but physical (e.g., books) [143]. For the shopping portal `amazon.com`, for example, the dominating costs comprise building and maintaining the physical storehouse and the web portal (which includes web design, back-end databases, security measures, credit-card clearing, and so on). These costs are almost independent of the number of customers surfing the site or buying items, most definitely though less correlated to the number of customers than in a traditional (book)store. The remaining variable costs consist of finding the items in the physical warehouse and shipping them to the customer.

As established above, though, vendors try to exploit every customer's readiness to pay, with the condition that variable costs are covered [121]. If for an increasing part of both products and of the processes to deliver them, the variable costs decrease, the tendency towards an increased usage of price discrimination is easily understood.

6.1.2 Product Expiration

A different problem encountered by stores, and especially by supermarkets, is that of depreciated goods. Such goods have to be thrown away because they have excessed their date of expiry. This problem amounts to roughly 0.6% of turnover for retailers in the food sector [136]. What may not seem much at first glance, becomes more impressive when looking at some statistics of food sales. According to the Office of National Statistics of the UK, the average weekly sales in "predominantly food" stores amounted to 1.7 billion pounds in 2000, with an increase of 22%

(to 2.08 billion pounds per week) in just the six years to 2006 [120]. This amounts to a turnover of roughly 161 billion Euro per year for Britain. The US Food Marketing Institute estimates the total supermarket sales 2006 in the US to 499.5 billion US dollars (or 374.75 billion Euro) [82].

Not only does this mean hundreds of millions up to some billion Euro value of depreciated goods per country and year. When putting the 0.6% in perspective to the small profit margins in the food retail business (1.46% net profit for the US according to the same Food Marketing Institute [82]), the problem's extent becomes clear.

This is why [136] argues that electronic means should be used to monitor the expiry date of goods in shelves and in the warehouse. By offering progressive discounts when goods approach their best-before-date, and thus avoiding total loss of value, an important part of this sum could be saved. In 2000 though, only 9% of US stores were monitoring the depreciation of their goods electronically, while, for example, electronic customer profiling – according to the authors a much less effective method for increasing profits – was used by 15% of stores.

6.2 The Ladenburg Shopping-Scenario[2]

The Ladenburg future shopping scenario includes several new or improved business processes enabled by ubiquitous computing technology, as mentioned at the outset of this chapter. Among them are to be found: automatic checkout enabled by technology such as RFID tags embedded in the products as well as RFID gates at the store's exit; a shopping assistant included in the shopping basket, which downloads the shopping list from the customer's PDA, continuously checks which products are still missing, and verifies each product's ingredients for possible conflicts with allergies in the customer's household; a quality and price assistant who downloads product reviews for non-food products (such as DVD-players) entering the basket as well as the prices from other physical and online stores; or physical navigation through the market along an optimized route towards the items on the shopping list.

However, the scenario's main focus clearly lies on the concept of indi-

---

[2] The term "Ladenburg shopping-scenario" does not only strictly comprise the scenario from the scenario collection [30], but also the ideas spread throughout the papers emerging from the collegium, which will be referenced where appropriate.

vidual, dynamic pricing, seen in the scenario as an economic tool used to reach a better price discrimination among customers, as well as for tackling the problem of expired goods. This tool could soon be made possible by ubiquitous computing technology. Given the large interest of vendors to reduce the losses caused by product depreciation, and also their interest to increase profits by using more numerous and more subtle means of price discrimination, the economical tool is likely to be used as soon as technology makes it available.

By pursuing such well-defined goals (reducing good depreciation, allowing for more price discrimination), the shopping scenario is clearly a problem-oriented scenario.

6.2.1 Individual, Dynamic Pricing

Individual, dynamic pricing can be seen as a step beyond nowadays' price discrimination. To sell similar goods to distinct customer segments for different prices (i.e., to realize third-order price discrimination as described in section 6.1), vendors can differentiate prices regionally or seasonally. The regional differences can range from large, country- or continent-sized (e.g., the above-cited example of the European car-market [158]) to rather small geographical differences, where a few dozens of meters between a shop inside the train station (but conveniently placed on the people's way to the platforms) and a shop outside the station can account for a large price difference for soft drinks, for example. Seasonal fluctuations in prices can also vary from large to relatively short timescales. There are some typical instances of prices largely depending on the season, as the examples of ice-creams being cheaper in winter than summer or of the end-of-season sales of clothing products show. There are, on the other hand, also goods that change their price as often as from day- to night-time. Many telephone companies, for example, charge a lower connection amount for calls at night (when the demand is lower) than during the day. Electrical power is also often cheaper at night, although the production costs for most types of power plants are the same for night and day power.

All these examples, however, although being characterized by fluctuating prices, rely on rather simple pricing models. Furthermore, prices are not individual. Every household being under contract with a particular power supplier will pay less during the night. The Ladenburg shopping scenario envisions that through the use of ubiquitous computing technology, price discrimination could be taken to an entirely new level of granularity, bringing it closer than ever before to the up to now merely theoretical concept of first-order price discrimination. Prices for goods would not only change with yet unseen dynamics (possibly on a per-minute-basis) according to the context they are in, every customer would also pay his or her individualized price depending on the shopping history.

This idea is underlined in the scenario by a future supermarket, where all products, including – as a deliberately provocative example – milk bottles, are equipped with sensing, computational, and communication facilities. Bottles communicate with each other and the market's computing infrastructure, so each of them knows how many other bottles are on the shelf, whether there is more supply in the warehouse, and also the expiration dates of all other bottles. The shelf has all this information as well plus other information influencing milk demand – like time, season, or weather outside. Plus, of course, the overall history of milk buying over the past months or years. Dependent on all these parameters, milk bottles set their prices dynamically. For example, while approaching the expiration date, a bottle will decrease its price, so customers will be tempted to buy it and not grasp for another, fresher one. Same would happen on the third rainy day in sequence, when sales decrease and part of the stock risks to depreciate. On the other hand, when sales are exceeding expectations, bottles notice they become lonelier in the shelf with time passing by. If the warehouse is also empty, remaining bottles will steadily increase their prices, as long as people are buying. Regular shop customers, however, are still shown the lower price as an acknowledgement of their fidelity. Thus, prices are highly dependent on the context as well as on customer's behavior. [98]

Price discrimination (the vendor setting prices largely independent of production costs, but according to the demand and the prices buyers are prepared to pay) is not, however, the only way that prices could be influenced through an extensive use of ubiquitous computing. The products' taxes, which obviously influence their end-consumer price, could also be assigned with a yet-unknown precision. As we have argued in [16], if the history of products is known (i.e., where they have been produced, their transportation means and the route taken), more precise and finer macro-economic control will be possible. Goods could for example autonomously determine their ecotax rate using the means by which they have been transported from their production location to the sales point (i.e., truck vs. railroad). Taxes could also vary depending on the length of the journey, in order to favor regional producers. Other product characteristics would also depend on the product production history. For example, milk could automatically be classified as organic if its history shows that it comes from a suitably accredited source.

6.2.2 A First Analysis of the Scenario

In a slight deviation from the process of scenario-driven prototyping as defined in the last chapter (which would have required at this point a prototype to underline some of the scenario's main features), we conducted a first interdisciplinary analysis at this early stage.[3] Having already had a suspicion that – although economically meaningful and in the near future technologically possible – such continuous price variations for the simplest everyday items might be bothersome and hit a wall of non-acceptance, we set out to analyze the scenario together with economists and philosophers of technology.

And indeed, Prof. Christoph Hubig, technology philosopher at the University of Stuttgart, Germany, and member of the Ladenburg collegium's outer circle, pointed out that a large obstacle for the acceptance of a particular technology appears when the technology in question implies the loss of value for past experiences [67]. When the experiences previously gathered by people deprecate at a fast pace, there is no knowledge sustainability. As we have thus argued in [16], "people do not feel at ease in highly dynamic environments, and there is a lot of data and information in everyday life that remains valid for quite a long period of time, e.g. food prices in our favorite supermarket, or public transportation prices. It is the sustainability of information that permits us to use acquired knowledge and prior experiences to cope with future situations and tasks. This raises the question of how far people can still cope in an over-reactive world of ubiquitous computers that has lost an element of inertia. In a highly dynamic world, the sustainability of knowledge risks being lost. An experience that was valid and useful one minute could become obsolete and unusable the

---

[3] This seeming slip, however, merely represents a deviation from the formalized steps of the scenario-driven prototyping process, which are – an inherent property of any formalization – rigid. The deviation, as will shortly be shown, very much complies to the *spirit* of scenario-driven prototyping though, i.e., not to waste resources on "lunar" technology [106] (and thus to draw irrelevant or even misleading conclusions), but to focus the analysis on meaningful, realistic, likely high-impact technologies.

next. A move towards highly dynamic systems could, therefore, have serious implications, such as the loss or the accelerated devaluation of long-term experiences, which could, in the long term, contribute to an increased uncertainty and lack of direction for people in society." In the more restricted area of supermarket shopping, this danger can be clearly seen: What could the benefit of past shopping experiences for the weekly family shopping still be, if prices varied every minute, and the price range would be wide open?

Another problem that has been pointed out by Prof. Hubig, and which could also lead to acceptance issues of ubiquitous computing technology, is the possible lack of usage traces [81]. People need feedback when they use technology, which helps them understand both the state of the system, and whether an intended action has been accomplished or not. By picking the handset of a telephone, for example, it is immediately evident whether the phone works correctly or not through the presence or absence of the dial tone. For the same reason, when activating the turn signal in today's cars, a typical ticking noise can be heard. The noise had originally been a side effect of the electrical switches used to open and close the electrical circuit of the turning lights. Nowadays, the switches being electronic, they could work soundlessly. However, the ticking sound is still artificially produced to provide drivers with a feedback for their action.

If a technology does not provide such feedback, a whole range of insecurities will arise (e.g., in the form of uselessly repeating actions that have already been executed, vainly trying to perform an action via a system that broke down unnoticed, or giving up actions in the first place if their result is impossible to tell), which will in turn lead to frustrations and rejection of the technology. In the context of the supermarket with highly dynamic prices, such insecurities could easily arise. Even if they would accept the loss of past shopping experiences and embrace prices continuously changing on the shelf in front of them, how can customers for example be sure that the price they accept at the shelf will be the one they will pay at checkout?

Finally, the problem of societal fairness has also been pointed out. In the first version of the shopping scenario, regular customers or such with a high income, are provided with special prices, offers, or independent tests of the products they buy. If manufacturers and traders are able to precisely evaluate the consumption behavior of their customers and, leading on from this, present individual customers with individual offers, a new quality of digital divide would quickly arise, one that extends into various branches of the real world. David Lyon, Professor of Sociology at Queen's University in Canada, calls this process "social sorting" – "Categorizing persons and groups in ways that appear to be accurate and scientific, but which in many ways accentuate difference and reinforce existing inequalities" [107].

These problems spotted within the scenario are not of mere theoretical nature. The example of the online storehouse amazon that tried to charge individual customers different prices for DVDs [141], but had to quickly take back this practice due to massive customer criticism, shows that customers are not willing to accept neither a decrease in the value of their shopping experience, nor the feeling that other, "more equal" customers (as George Orwell would put it) pay less for the same everyday item.

6.2.3 New Scenario: Individual Accounting of Traffic Costs

After these first interdisciplinary discussions, it became clear that while the idea of individual, dynamic pricing has a large potential within a ubiquitous computing environment, it had been originally presented in a less suited context. Thus, we conducted further discussions with economists, to see whether we can find a better application context to stress out the concept. Thereby, we understood that a segment of goods, rather different than the originally considered supermarket items, exists, where the idea of individual and dynamic pricing is likely to take off – and that in order to provide the scenario the necessary realism and impact, we should use such goods to present the concept. The goods for the new scenario should differ from supermarket products or DVDs in three main aspects: "First, they should be expensive enough so that the deployment of the necessary ubiquitous computing infrastructure does not increase their price significantly. Second, both sellers and buyers should benefit – or at least a fraction of the consumers that is large enough to represent an interesting market should gain from such augmented products. Third, charging differently for the good should not be uncommon nowadays – the stronger and more noticeable today's price discrimination for the good, the more likely consumers will accept new types of price discrimination in the same area" [29]. Especially this third point would make the whole difference in terms of acceptance from customer's side. In an area where customers are already used to pay different prices for the same good (especially if they also understand and accept the reasons for the differing prices), a further discrimination is quite obviously more likely to occur.

After being presented with the concepts of ubiquitous computing, its possibilities, and the problems pointed out for our original scenario, Dr. Jochen Jagob, economist at the Technical University of Darmstadt, has been the one to suggest insurances as a domain much better suited to present and analyze our concept within. In terms of acceptance, it is easily seen why insurances are a well-suited domain. It is already widely accepted nowadays that insurance companies try to understand the customer as precisely as possible in order to accurately assess his or her risk, and offer an insurance premium proportional to that risk. Car insurances, for example, are largely based on some parameters that have been proven to be relevant for the accident risk, such as age and experience of the driver, his or her accident history, or the type and power of the car – criteria that do not get questioned by customers. It thus seems that having even more personalized and – as a new aspect – also dynamic and behavior-dependent insurance premiums is more likely to be accepted by customers than, say, price-changing milk bottles.

However, the probable acceptance by itself does not yet imply that insurances represent a likely application domain for individual and dynamic pricing. The other condition that has to be fulfilled – even before thinking of acceptance – is that there is an interest from insurers to take up a new quality of discrimination between their customers. In the motivational part of this chapter we have given two arguments for the shopping scenario in the context of a supermarket: the interest of vendors to pursue price discrimination, and the depreciating products in warehouses. While the second argument obviously cannot apply to insurances, we will argue in this section that discriminating individual customers (instead of merely dividing them into customer segments, as done today) not only makes sense for the insurance market, but that it would have as direct consequence the tackling one of the largest problems insurers face nowadays – the problem of information asymmetry.

Information Asymmetry

Information Asymmetry has been defined back in 1970 by Berkeley professor of economics George Akerlof in the seminal article “The Market for Lemons: Quality Uncertainty and the Market Mechanism" [4]. Over the next decades, the concept has proven to be so relevant for insurance markets that Akerlof received the 2001 Nobel Prize in Economics for his work. In his article, he recognized that when one market side is better informed about the traded good than the other, the resulting unbalance of the marketplace leads to several negative effects, up to the point where some markets might fail to work.

To illustrate the phenomenon, Akerlof looks at the well-known example of the used vehicle market. In this market, the seller has more information about the car's condition than the buyer. Not knowing the history of the vehicle, the buyer is suspicious: will he buy a well-maintained, high quality car; or one with accidents in the past, leaking oil, and with an almost broke-down engine – in other words a 'lemon'? As we have argued in [32], since the buyer is suspicious, he will always try to pay the price as if the vehicle was indeed a lemon. However, for such a low price nobody will be willing to sell a high quality car. Thus, without further measures to break the information asymmetry (which exist and will be addressed shortly), on the marketplace either an in-between price would be established (disadvantageous for sellers of good cars and buyers of lemons), or – if buyers will only be willing to pay the lemon price – only lemons will be sold. In the latter case, high-quality cars would be downright pushed out of the marketplace by lemons.

Obviously though, high-quality used cars are being sold on the market. In order for this to happen, Akerlof argues, the information asymmetry (or at least a part of it) has to be broken. For the example of used cars, means of credibly proving the quality of the car have been established over the years, such as presenting a full-service history certified by a recognized garage, having a recently dated thorough inspection, or providing the buyer with a warranty.

In insurance markets, from the outset the area targeted by Akerlof's theory, it is – at first glance maybe surprisingly – not the buyer who has less information about the traded good, but the seller. Only the insurance buyer knows precisely the state of the traded good – may it be his health, his car, or any other good that is being insured. This information asymmetry results in two negative effects, which are specific for insurance markets: adverse selection and moral hazard.

Due to the incomplete information about the behavior of each of his customers, the insurer will proceed and divide them into classes of risk, according to his best knowledge of their individual risk profile. However, although better than no criteria at all, these customer segments are often only roughly sketched, being based on the few criteria the insurer is able to determine. As a consequence, inside each customer class (all paying the same insurance rate) there still exists a large spread of the individual risks.[4] *Adverse selection* denotes the cross-financing thus taking place within every risk class – from the lower risks (who pay more than they should for their actual risk) to the higher risk customers (who pay less). *Moral hazard*, rooted in the insurer's lack of information about the customer's behavior as well, denotes the customer's tendency to handle goods less careful once they've been insured, since the lion's share of possibly occurring damage costs has been transferred to the insurance company.

Since the increased costs due to moral hazard are supported directly by insurers, it is easily seen why they have an interest in reducing the moral hazard. As for the reduction of adverse selection, not only the customers paying too much today would like to see it reduced, it is also in the insurer's best interest to do so. As Oberholzer [119] argues, this would gain them market shares in the attractive segment of low risks (through the lower rates the insurer would be able to offer them), as well as help insurers to get rid of the bad risks (who – through the same mechanism – would have to pay higher rates and would thus be likely to change the insurer).

As in the case of the used vehicles market, means to reduce the information asymmetry for insurance markets already exist. Numerous insurance contracts, for example, leave some of the financial responsibility with the customer (via franchises), aiming at reducing the moral hazard. Furthermore, inspectors might examine and reconstruct the course of events to see whether the insured has committed gross negligence, case in which he would lose the right for damage claims. On the other hand, insurers often offer several types of insurance contracts, some more restrictive but with smaller premiums, others more generous and according higher premiums, aiming at a "self selection" [50] of customers into higher- and lower-risk classes. However, all these measures can only reduce the problem's magnitude; insurance markets are still largely characterized by adverse selection and moral hazard [50].

---

[4] For a discussion on the imperfect categorization in insurance markets, see [79].

Information Asymmetry for Vehicle Insurances

Among the different existing insurance markets, vehicle insurances seem particularly likely to adopt technological means that would enable insurers to offer substantially more individualized premiums. There are several arguments to sustain this. First, the three criteria put forward at the beginning of section 6.2.3 seem to be all achieved: Vehicle insurances (and even more so the vehicles themselves) are expensive goods, so that the costs for the infrastructure needed to give the insurer more information about the customer's behavior (i.e., individual driving style) would account relatively low.[5] Second, as argued above, there is an interest for more personalized insurances on both sides of the market: Insurers would reduce adverse selection and moral hazard, and the below-average customers in today's risk segments would save money. Third, there already is a strong price differentiation between individual customers nowadays, which leads to a simpler acceptance.

However, there are also other, more subtle arguments to favor vehicle insurances over other type of insurances when considering individualized, behavior-dependent pricing. As we have noted in [28], "there are several other examples of costs that could easily be allocated to their originators, yet they are burdened by the society as a whole. One such example are health insurances. Instead of evaluating the individual risk of illness based on age, gender, and health history, many countries have decided to spread those costs throughout the society, willingly cross financing the elderly or the ones with chronical diseases from the young and healthy." This touches an often-encountered problem of ubiquitous computing visions and scenarios, as we have pointed out in the previous chapters: technological possibility is not linked in any way to economic or societal meaning; but all these criteria are needed for an application to be realistic. The economic drivers for individualized, behavior-dependent vehicle insurances have already been presented. Societal value will be discussed below; however, it already stands to reason that vehicle insurances don't have such a high societal value as, for example, health insurances. The communitarian feeling of stepping in for the "weak" (i.e., risky drivers) by sharing their costs is not as pronounced as in other cases (when truly weak members of society such as chronically ill need our financial support), making thus

---

[5] Moreover, as will be argued later in this chapter, much of the infrastructure needed is already present in vehicles for other purposes, such as traffic safety, theft protection, or satellite navigation.

highly individualized vehicle insurances seem more realistic.

A further aspect is the fact that nowadays' vehicle insurances fail to account for parameters that largely influence the accident risk. Besides the unfairness of such system, other societal drawbacks might occur. Litman [102], who proposes mileage-based vehicle insurances, even argues that today's rigid insurance premiums are both economically and ecologically obsolete: "What would be the consequences if gasoline were sold like vehicle insurance? With gasoline sold by the car-year, vehicle owners would make one annual advance payment which allows them to draw gasoline unrestricted at a company's fuel stations. Prices would be based on the average cost of supplying gasoline to similar motorists. Unmetered fuel would cause a spiral of increased fuel consumption, mileage, and overall vehicle costs, including externalities such as accident risk, congestion and pollution." Instead, the insurance should be related to the mileage driven because, all other parameters equal, there is a strong correlation between driven distance and accident risk. Connecting the insurance rate to the annual mileage would be fairer and economically more sensible. Moreover, since a larger fraction of vehicle costs would depend on the driven distance, it would also have a positive environmental side-effect. Oberholzer goes even further and builds up a detailed matrix of how much the insurance kilometer should cost depending on two factors: type of road (highway vs country road vs city) and the hour of driving [119].

Finally, more individualized vehicle insurance premiums seem to be not only economically and ecologically, but also societally meaningful. As we have argued in [34], "today's car insurance schemes typically divide drivers into about two dozen different risk categories, using only a few criteria such as the driver's age, gender, driving experience, place of residence, or car model. While all these parameters are being determined before the insurance goes into effect, the actual behavior of the driver after signing the policy (e.g., a safe driving style) will reflect only slowly on his or her insurance rate, typically only after one or more years of accident-free driving." Young drivers with a safe driving style, start – like all young drivers – with the highest premiums, only being able to reduce them over several years of accident-free driving. During these years in which they prove themselves to be the safe drivers they have always been, they get punished for the fact that of incidentally belonging to an – on average – high-risk group.

Road Pricing

Highly individualized vehicle insurances seem to be meaningful from fairness, business, economic, ecologic, and societal perspectives, and appear thus as a well-suited domain for realistically assessing the impact of individual and dynamic pricing. As opposed to the supermarket products presented in the original scenario (where the price fluctuations were purely demand-driven with the aim of maximizing the vendor's profit), the individual pricing for vehicle insurances would directly reflect the insured risk. Thus, the envisioned system offers an individual accounting of traffic insurance costs according to the originator principle.

However, there is also another important class of traffic costs which are supported nowadays by the society as a whole, but could be accounted for by their originators. These include costs such as air and noise pollution costs, as well as the economic losses (time and gas consumption) which occur due to traffic jams. Over the last years, the public debate to account such costs to their originators (especially with the hope to see them reduced as a result) has increased in intensity. Most proposals suggest to substitute the (flat) vehicle tax with a usage-based road pricing.

Road Pricing is a tool for regulating the traffic flow through selectively penalizing the driving on specific roads at particular times or under specific conditions. Deploying a road pricing scheme may have several political or societal aims. For overcrowded city centers, it may be deployed for replacing the regulation of traffic through queuing (the "communist" solution) by a free-market mechanism. It may also be used to steer the traffic away from some streets to others (by penalizing the former more) or to other means of transport. Road pricing may further pursue environmental aims, like reducing emissions or noise levels. Finally, it may simply be used to raise money for the maintenance of the road infrastructure. However, as [47] argues, whatever the main reason – financing, improving the environment, or managing traffic and improving accessibility – a road pricing system will have all these effects to a certain extent.

Road pricing has become increasingly popular over the last years, since the two traditional tools for charging drivers, fuel and vehicle tax, are rather coarse and cannot fulfill all of the above mentioned aims. The annual vehicle tax penalizes people for owning dirty cars, but this says nothing about the actual pollution caused by those cars. Since the car's overall pollution is the product of its emissions per distance unit and the car's usage, this flat tax fails short of fulfilling its environmental aim. Fuel tax is better at penalizing people for the consumption of gasoline, but does not look at the other side – how dirty the emissions resulted from that consumption really are or under which circumstances the gasoline has been burned (e.g., ozone levels). Moreover, neither tax can have a traffic management effect [148].

Hence, several places worldwide have started to deploy road pricing systems as a complementary tool to the existing taxes. There is a wide spread in the level of detail that existing road pricing systems take into account. At one end are rather coarse systems, such as the London Congestion Charge. There, motorists are charged once every 24 hours for the permission to drive in the city center [47], regardless of the actual usage during these 24 hours. Even so, the introduction of the Congestion Charge has reduced traffic by 15% and increased speed by 22% in central London [148]. At the other extreme, having a much more detailed usage model, stands the ERPS (Electronic Road Pricing Scheme) from Singapore. There are different taxes for the usage of distinct roads, and they also vary with the hour of driving. Furthermore, every three months, the whole price structure is analyzed and readjusted [47].

Many other places play with the idea of road pricing systems: Viennese officials are thinking of a city-wide street usage charge of 2 to 8 Euro-Cents per kilometer [146]. A large study has been carried out by the Swiss Center for Technology Assessment to investigate the public's acceptance of a generalized road pricing scheme, envisioned by parts of the government [127]. Britain also thinks of a nationwide, satellite-based road pricing system, no earlier than 2014 though [147].

New Scenario: Pay-per-Use / Pay-per-Risk Accounting of Traffic Costs

Thus, the new scenario, as described in [28, 34, 35], envisions a ubiquitous computing system installed in vehicles, which continuously analyzes driving behavior and driving conditions. Depending on these, and according to specific rules, the system is able to compute an arbitrary number of different types of traffic costs and account them to the motorists causing them, according to the originator principle.

Some attributes of the scenario (and, subsequently, of the proposed system) are worthwhile noticing. From the outset, an important novelty of our proposal has been to have an open platform that allows for a discretionary number of types of traffic costs to be computed and accounted with. Two such types of costs have already been presented – the vehicle insurance and the vehicle tax in form of road pricing. But virtually any kind of traffic cost can be included. To be thought-provoking and show the power of such technology, in the prototype that we have implemented (which will be presented in the next section), one of the considered "costs" are traffic fines. These are automatically issued and accounted to their "originator" when the car surpasses the speed limit, while at the same time the police is informed.

While – as the related work in section 6.5 will show – some road pricing systems and behavior-dependent insurance prototypes have been proposed before (albeit usually a lot simpler than the system we have developed), all proposals so far have been dedicated systems, for one type of traffic costs only. The power of an open platform, though, is more than the sum of its (initial) parts, as the example of the police tracking of "speeding costs" shows. Our scenario has been the first to propose such generic approach on a theoretical level and the prototype developed the first one to prototypically implement this idea.

One more attribute that became clear while developing scenario and prototype is the necessary continuity of the cost function. Only by having a continuous cost function, the complexity and fine-granularity needed for the correct assessment of the individual costs of millions of drivers can be guaranteed.

A necessary first step for the assessment of costs is the ability to express behavior into costs. What is the precise risk of an accident, for example, when taking a curve of a certain radius with a certain speed (and corresponding cross acceleration), given the prevailing traffic? What are the expected average damage costs to the type of car that is being driven if an accident does happen in that curve? The product of these two factors represents the average expected damage sum for the insurer. What are further the health and environmental costs of one kilogram of emitted carbon dioxide, for the specific environmental conditions prevailing, such as temperature and ozone level? Obviously, monetarily expressing some of the costs caused by traffic, such as the environmental impact or health damages inflicted upon others, are not easy tasks. There is an entire research domain concerned with this task,[6] which has not been able to provide a final answer yet,

---

[6] See, for example, [48, 91].

but which has several promising approaches. A detailed discussion lies outside the scope of this thesis, as it has not been part of the individual and dynamic pricing scenario. Including the findings from this area is nevertheless of great importance for realistic and fair pricing schemes.

However, for being able to charge the exact costs to their originators (after defining such functions which transform behavior and prevailing conditions into costs), it stands to reason that a necessary second step is to continuously assess these costs. Both environmental conditions and driving behavior change at a rapid pace, so that the precise costs caused by the motorists are inherently momentary in nature and have to be continuously measured. Mathematically speaking, the costs have to be expressed as currency unit per time or distance unit, while the aggregated costs will be the integral (over time or distance, respectively) of this continuous function. Some of the costs – which occur while the car is not moving – could not be expressed as currency over distance and have to be expressed as currency over time. Using only time would even suffice, when the vehicle is moving any currency-over-time cost can be easily transformed into currency-over-distance, since the distance is time times velocity (an attribute that we can assume to always be available, since it has always been measured in vehicles). However, since costs per kilometer of driving is a more familiar concept than costs per minute of driving [119], the scenario and our later prototype allow both, using whenever possible (e.g., for insurances), the currency-per-distance concept. Moreover, to allow for even more flexibility, the scenario also requires that singular monetary events should be allowed. Such property is used, for example, for the speeding fines which are by nature not continuously occurring costs, but one-time-events.

A final attribute of the system envisioned within the scenario is providing the driver with real-time feedback about how his or her driving style influences the different types of monitored costs. Thus, the driver would have indicators for the momentarily insurance and road pricing costs, for example, in a similar way that some cars show the current gas consumption. The hope is that, by providing a direct feedback of how dangerous (and, in turn, how costly) some of the maneuvers are, this will encourage drivers towards a more defensive and safe driving style.

Looking back, what used to be a supermarket shopping scenario in which products changed their prices to avoid depreciation, and make use of as much as possible from each customer's readiness to pay – a scenario describing a technologically challenging, economically questionable (due to the relative high costs of the system compared to the value of the products regarded), and above all with unlikely customer acceptance application – became a scenario describing a more mature, economically relevant, technologically realistic system, which could exhibit profound societal implications and where important trade-offs have to be societally agreed upon before the system is deployed – a system for the precise, fine-granular, and individual accounting of traffic costs to their originators on a pay-per-use / pay-per-risk-principle.

6.3 The Prototype: The Smart Tachograph

Following the scenario-driven prototyping principle (as presented in Fig. 5.2), at this point the second step had to be done – developing a prototype emphasizing the main features of the scenario, in order to further and deeper investigate the implementation requirements and the subsequent economic and societal consequences of such precise, individual, comprehensive, and continuous monitoring of driving behavior and subsequent cost allocation. To do so, we have proposed and developed the *Smart Tachograph* generic platform [28, 34].[7] This section presents the main aspects of the Smart Tachograph: the concepts that have been implemented, and it's general hardware and software design.

6.3.1 System Features

As discussed in the previous section, the Smart Tachograph's main objective is to provide an open, flexible platform for the individual accounting of various types of traffic costs, according to the originator principle. To attain this main goal, two conditions have to be fulfilled. Firstly, the system has to allow for different types of traffic costs to be computed in parallel. Since the system is conceived as a generic platform for all types of traffic costs, however, the system should ideally also be easily extendible for new types of traffic costs. Moreover, already existing cost types should be easily modifiable – since we are talking about such a fine-granular calculation encompassing numerous attributes, it is likely that changes in the rules for computing costs will occur over time.

---

[7] The software of the Smart Tachograph has been developed together with Christoph Plüss, as part of his master's thesis [125]. Most drawings in section 6.3 are adapted from his master's thesis.

Secondly, since all these costs depend on momentary environmental data and on the driver's behavior – both of which are measured by sensors – the infrastructure should allow for the continuous monitoring of different kinds of sensors. As the types of traffic costs might change over time or new sensorial data might become relevant for already existing types of costs, the needed sensorial input might change as well. Adding new sensors to the system should therefore be an easy task, similar to the way new types of traffic costs should be easily added.

Overall, the Smart Tachograph's main aim is schematically presented in Fig. 6.1. It allows for different types of traffic costs to be accounted to individual motorists, by transforming sensory input into costs using predefined rules.

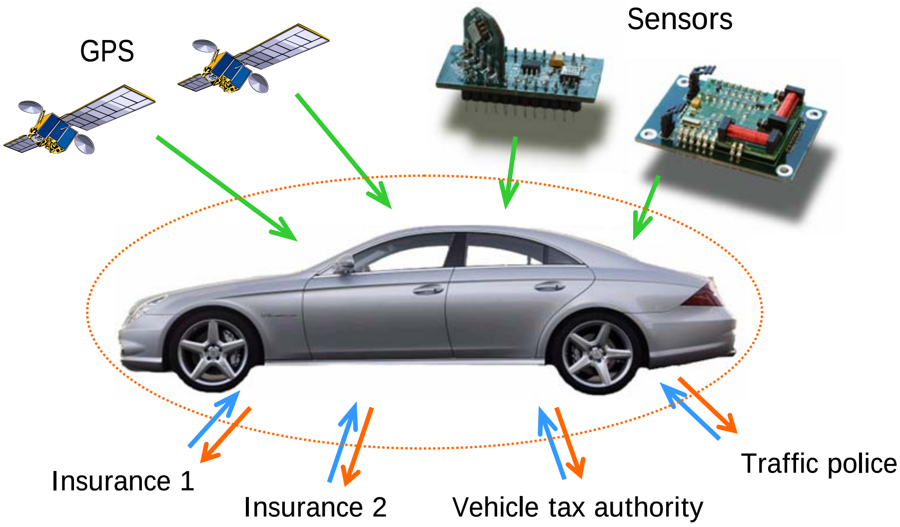

Figure 6.1: Schematical view of the Smart Tachograph, which allows for different types of traffic costs to be billed to the motorist according to predefined rules and sensorial input.

We will present now these main concepts in greater detail, as well as some other concepts that have been included in the Smart Tachograph in order to allow the analysis of further significant technological, economic and societal implications.

Accounting Authorities for Various Types of Traffic Costs

To denominate an entity who has an interest and the right to charge motorists some sort of traffic costs, we use the term *accounting authority*. For any type of traffic costs that can be accounted via the Smart Tachograph, there is one sort of accounting authority. The three types that have been included so far in our prototype are:

- **Tax authority**. A tax authority will charge motorists a road pricing tax, with the aim of regulating traffic, steer traffic away for some streets to others or to entirely different means of transportation, make motorists pay for the environmental and economic damages they inflict to others by pollution and creation of jams, or for a combination of these aims, as described in the scenario.
- **Insurance product**. Insurers charge the pay-per-risk insurance. In the real world, an insurance company will typically offer different types of products which cover semantically different types of traffic costs, such as liability insurance or own damage claim. For the sake of simplicity, we use the paradigm as described before – for every type of insurance product there is one type of accounting authority. For "liability insurance" and "own damage claim," although they might be offered by the same insurance company X (or by several companies A, B, C, etc), there will be two different types of accounting authorities.
- **Police**. The police can also charge motorists a certain type of traffic "costs," namely fines.

Required and Non-required Authorities

For some accounting authorities, there can be one instance only. In our prototype, for example, the "police" and "tax authority" classes of accounting authorities are only instantiated once, as it would most likely happen in real-life. For other types of accounting authorities, as for vehicle insurances in our prototype, several parallel instances offer their services, the motorists having the choice.

However, for some types of accounting authorities, although motorists do have the choice, one instance of the corresponding accounting authority is required. In our prototype, the "liability insurance" is such an authority – customers may choose between many products from different companies, they are nevertheless required to have chosen one.[8] Other accounting authorities are entirely optional. In our prototype, the "own-damage-claim" insurance is such an example.

To reflect these three classes of accounting authorities, the Smart Tachograph requires that any type of accounting authority be registered

---

[8] In order to make the prototype reflect reality as close as possible, we have adopted here the general policy of European legislation, which requires motorists to be insured against damage inflicted to others.

in one of the following three lists:

- **Compulsory**. This list contains the types of accounting authorities for which a singular, required instantiation exists. The billing mechanism of each of these singular instances must be active at all times.
- **Required-select**. This second list also contains mandatory accounting authorities. However, since motorists may choose among several offers, the system only checks that for each item in the list (i.e., each type of authority contained here), there is one active instantiation.
- **Optional**. This list contains optional types of accounting authorities. Customers might have instantiations of the items in this list on a voluntary basis or not.

System Flexibility

The world of highly dynamic and quickly adjusting prices, so ubiquitously present in the first shopping scenario, is still technologically possible for the new scenario of traffic costs. But, the main reason for the quickly adjusting pricing has shifted from pure profit maximizing in the case of the supermarket towards a fairer, individualized, economically meaningful billing of traffic costs. The vendor of the second scenario does not try to maximize profits by tapping the full potential of everyone's readiness to pay, but to have more gains by offering a more precise and fair business model.

Nevertheless, in a real deployment, traffic costs would most likely not only vary due to the driver's behavior and environmental conditions. This variation, the only one that we have talked about so far, is inherent to the nature of the system – every second, the costs produced and the risks taken by the driver are measured and get charged monetarily. However, the *calculation* of these momentarily costs follows a function which might be dynamic or static. We believe that in a real deployment the calculation rules themselves would be rather dynamic than static. Many circumstances are conceivable when the accounting authority might want to change the rules for the costs it charges to motorists: Due to new environmental laws, road pricing might get more expensive. Since the traffic reduction aims have not been achieved / have been exceeded, the road tax would get more expensive / cheaper.

Insurers could also want to change the basis of their calculations, because they have come up with new, more precise rules, for example, or because they want to make special offers in order to attract new customers or to reward loyal ones. Thus, the system allows for a quick and effective exchange of the basis of calculation. When an accounting authority decides to make such change, the system automatically notifies the drivers about the change. For the two classes of authorities where customers have a choice between several providers, they can accept the change or choose another, better suited offer.

Customers may, however, also change the provider at any other time they see fit. By supporting this feature, we have implemented the concept of *ride-by-ride insurance*, that we have introduced in [29]. Customers could then change their insurance on a regular, maybe even daily basis. Since insurers are likely to use different risk estimation strategies, they will probably offer different rates for different conditions. One company could offer the best rate on weekends, another insurance company would have the best price per kilometer on highways, and a third one could offer the best savings when driving at night in rainy conditions. As with telephony least-cost-routers, the car could autonomously choose the most favorable insurer for that day – for example, after the driver entered the destination in the car's navigation system. By supporting this feature, the Smart Tachograph allows for the analysis of interesting economic and political implications of ubiquitous computing technology, which will be presented towards the end of the chapter.

6.3.2 Overall Design

The overall system for the individual accounting of traffic costs according to the originator principle consists of the Smart Tachograph itself, which is a collection of hardware and software components placed inside the vehicle, but also of external systems, such as the software running on the computers of the accounting authorities and the protocols of communication between vehicles and these entities. This structure is presented in Fig. 6.2.

The following two sections will present the main component of our prototype, the Smart Tachograph itself, and the various external systems, respectively.

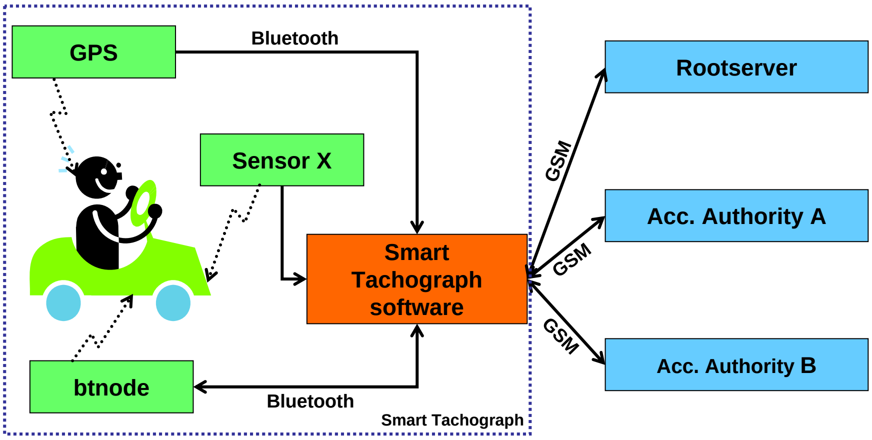

Figure 6.2: Top-level view of the system for the individual accounting of traffic costs including the Smart Tachograph and the other main involved instances.

6.3.3 Design – Smart Tachograph

The Smart Tachograph uses off-the-shelf ubiquitous computing technologies and a newly developed prototypical software infrastructure to allow for measurement of driving parameters, the transformation of those parameters into costs, and billing these costs to their originators. The software infrastructure of the prototype runs on a laptop computer that can be placed anywhere in the car. Any number of sensors can be attached to the system. The sensors gather data about the way and the circumstances in which the vehicle is being driven and send this information to the laptop computer. Several accounting authorities may evaluate this information and charge motorists on a pay-per-use basis. The software platform serves not only as a sink for sensor data and as a back-end connection for the accounting entities, but also as a front-end interface to the vehicle's driver.

Communication Technologies

The communication technologies used for the communication inside the vehicle and from the vehicle to external components of the system are presented in Fig. 6.2 as well. To guarantee an infrastructure coverage as large as possible for the communication between vehicle and accounting authorities, the GSM mobile telephony standard has been chosen. In future, GPRS or UMTS could be used instead. For the communication inside the vehicle from sensors to the laptop computer running the Smart Tachograph software, the Bluetooth wireless communication technology has been used.[9] In a real setting, the sensors could be hard-wired to one of the vehicle's computers, which would run then the traffic costs accounting software. However, this distinction is not relevant for the prototype. The Bluetooth technology, although inducing some pragmatical problems, has been chosen for convenience reasons.

Sensors

A small plastic box (Fig. 6.3) has been fitted for our prototype with a collection of sensors. It contains a GPS unit and a sensor board carrying two accelerometers (for longitudinal and cross acceleration), a temperature sensor, and a light sensor. Raw GPS coordinates are not the only information that can be obtained from the GPS unit. The current time is also encoded within the satellite signal, which allows us to infer the current speed as distance traveled over time. The data gathered by all these sensors is sent via Bluetooth to the computer running the Smart Tachograph software infrastructure. We used a Bluetooth-enabled GPS sensor, while the sensor board sends its data through a BTnode (in the lower part of picture 6.3). The BTnode[10] is a small computing device for sensor network applications equipped with Bluetooth communication capabilities.

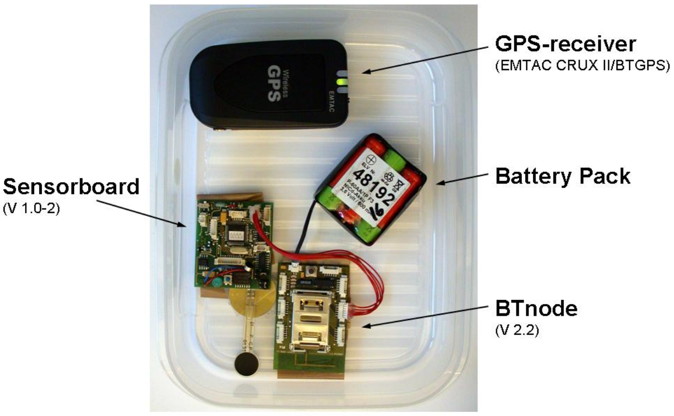

Figure 6.3: The plastic box containing the Smart Tachograph's sensors.

"Installing" the system is pretty straightforward. The box has to be placed in a spot where the GPS sensor can easily receive the satellite

---

[9] See www.bluetooth.com.

[10] See www.btnode.ethz.ch.

signals, for example underneath the car's windshield. The only other point to be ensured is that the sensor box is placed on an even surface and that it faces in the correct direction. Both conditions are needed for a correct functioning of the accelerometers. The controlling computer can be placed anywhere in the car, since it wirelessly communicates with the sensors.

We have chosen to deploy the mentioned type of sensors for two reasons. First, they all measure data that is potentially relevant to one of the envisioned accounting authorities. Having the raw GPS coordinates, the system can always determine on which street the vehicle is on and the speed limit for that street, using a commercial geospatial database installed on the computer. This information is obviously relevant for a road pricing scheme. But it could also influence the current accident risk, for example if the driven speed significantly exceeds the speed limit. Excessive longitudinal acceleration and especially excessive cross acceleration seem to be equally important indicators for a high accident risk. The temperature sensor would also offer valuable information (i.e., icy road), would it been placed outside the car. However, since it is placed in the sensor box inside the vehicle, its information is of no use for the prototype. Finally, the light intensity sensor gives information about light conditions, which could also influence the accident risk. The second reason to include these sensors in the Smart Tachograph prototype was that all of them are already available in a typical medium to high end state-of-the-art car. While being placed there for other reasons, they be reused in a real deployment for a system such as the Smart Tachograph. A modern car is equipped with acceleration sensors for the electronic stability programs, with a temperature sensor for signaling the driver a possible slippery road, and with light sensors for automatic headlight activation. Most cars today also come with a GPS system and navigation maps.

Apart from the sensors used in our example, many modern cars come with a variety of other sensors that could be used to determine the insurance rate or road pricing tax even more accurately. A distance sensor used in many cars as parking aid could be reused to measure the distance to the car in front. This information, correlated with the type of street and driven speed, is a major determinant for the current risks taken while driving (see section 6.4.3). If the car is connected to the Internet (e.g., via UMTS), it could also download environmental data that possibly determine its road pricing tax, like ozone levels or the concentration of carbon dioxide in a specific city.

Why haven't we used the car's sensors if they are already there and most likely more precise than ours? The practical reason was that we did not have access to the vehicle data bus, since the work has been academic research so far. The aim of the work lies not in the highest possible sensor precision, or in the most realistic approximation, but in creating a proof of concept for a generic traffic accounting platform and to analyze the various implications of deploying such a system. The fact that similar sensors already exist in vehicles only underlines the feasibility of the presented concept. Finally, and to our surprise, the information gathered from the sensors – especially the data from the acceleration sensors and the speed inferred from the GPS sensor – has proven to be surprisingly exact.

Smart Tachograph Software Architecture

The system has been developed in a Linux environment, since the BTnode used for sending the sensor data to the Smart Tachograph works best with this operating system. All the Smart Tachograph's software has been developed in Java, except of the BTnode software, where the C-environment is mandatory.

The main role of the software infrastructure is to query data from the sensors, and to mediate communication with the accounting places. Due to the flexible software design (see below), adding accounting entities is as easy a task as adding new sensors. At the time being, three different kind of accounting entities have been included in the system (see Figure 6.1): insurance companies, a vehicle tax authority, and the police. The traffic police has been included in order to show how powerful the paradigm of a smart tachograph is, and what far-reaching social consequences it could have.

The software architecture is divided into subsystems, which are depicted in Fig. 6.4. There is one core system and three auxiliary systems. The directions of the arrows in the figure indicate the flows of data between the subsystems. Configuration and initialization will be presented together with the external systems, since much of this process relates to connecting to the external entities. The other three subsystems are addressed below.

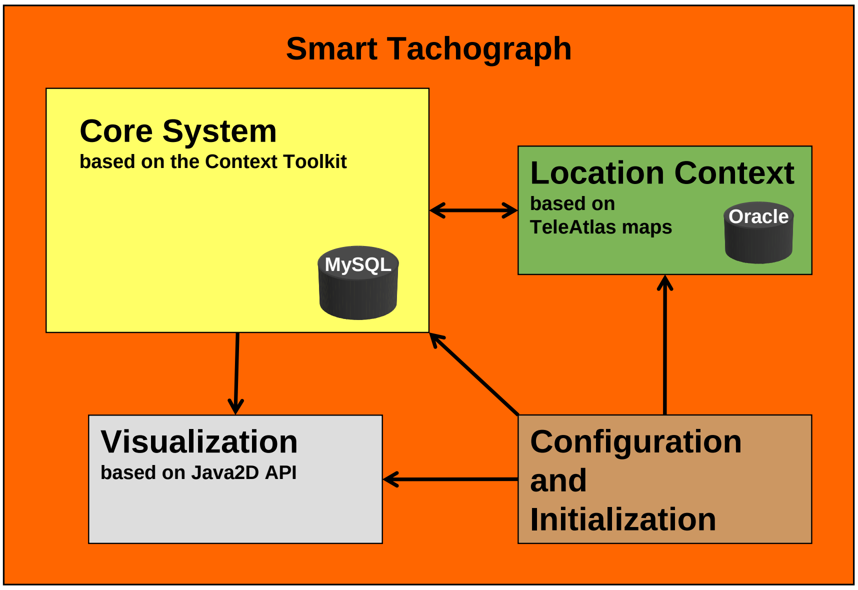

Figure 6.4: Top-level software architecture of the Smart Tachograph, divided into core and auxiliary systems.

Location Context

One of the most relevant information for the computation of traffic costs within the Smart Tachograph is the momentary location of the vehicle. Important, however, is not the raw position data supplied by the GPS-receiver, but the geographical or location context these imply (e.g., which sort of road or street the vehicle is being driven on, what is the maximum allowed speed on that street, etc).

Obviously, this kind of location context is already available in nowadays' navigation systems. Thus, the first idea has been to use a commercial navigation system and extract the location context from it. We have discovered, however, that vehicle navigation systems are purpose-built systems that use the internally stored maps only to present the driver the computed route to a destination point. We have found none to provide a programming API that would allow access to the location context computed from the raw GPS position.

Thus, we have eventually bought the needed maps from one of the only two worldwide providers of such detailed map material, TeleAtlas.[11] We have bought maps for the canton of Zurich, where we have later also tested the system. The road maps have been read into

---

[11] www.teleatlas.com.

an Oracle Spatial database[12] (an add-on to a core Oracle database, both of which have been installed on the laptop computer running the Smart Tachograph software). As Fig. 6.5 shows, a small software – `MapHandler` – has been written, which, upon receiving an SQL request from the core system with a GPS position, returns the type of street and the maximum allowed speed on that street.

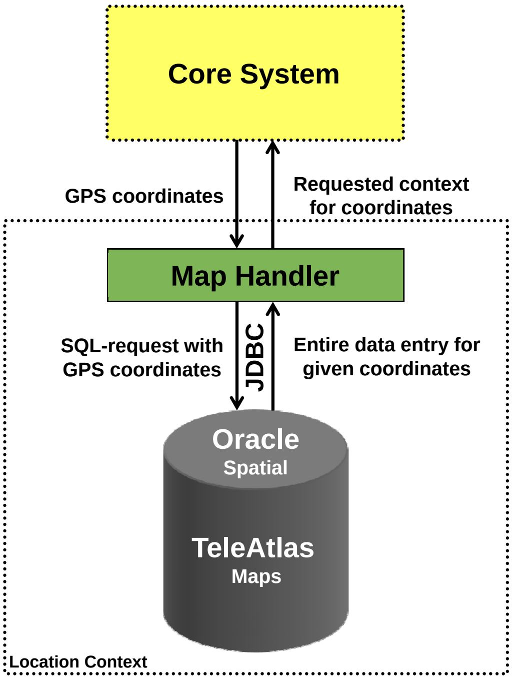

Figure 6.5: Schematic view of the location context system. The Map Handler's API replies with the needed location contextual data to a request containing the raw GPS coordinates.

The Core Software System

The core system accomplishes the main tasks of the Smart Tachograph. It establishes and maintains connections to the sensors, manages the accounting of the different traffic costs, keeps the archive of past sensory data, and manages the communication with systems external to the Smart Tachograph, such as the accounting authorities. These main tasks are depicted in Fig. 6.6.

To connect with the accounting authorities, the Smart Tachograph uses a plug-in-architecture as depicted in Figure 6.7. The system being built on top of Anind Dey's Context Toolkit [41] (briefly presented

---

[12] www.oracle.com/technology/products/spatial/.

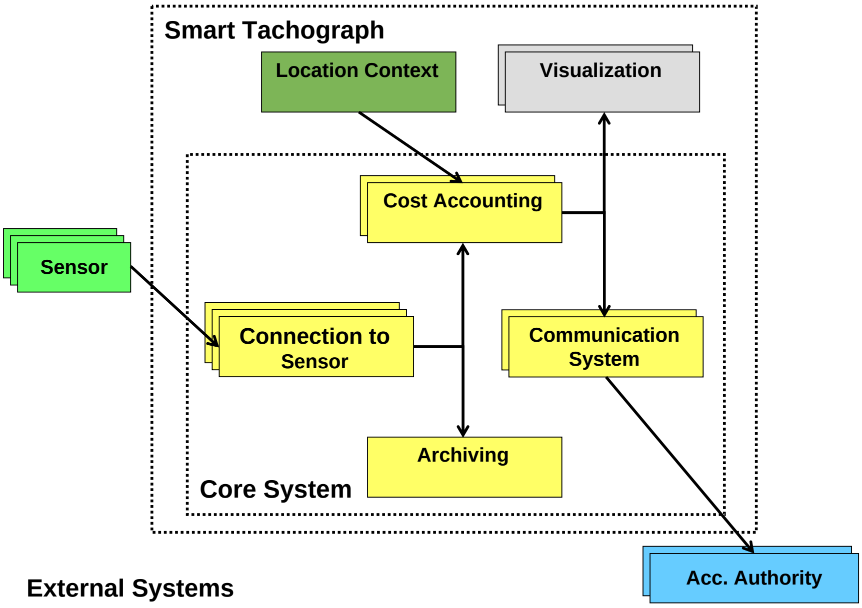

Figure 6.6: Simplified view of the Smart Tachograph's core software system with its main tasks.

below), every accounting authority is being represented by a Context Toolkit *server*. This server is registered through the Context Toolkit's publish/subscribe mechanism to receive all events from the GPS unit and the other sensors. The server sends this information to a Context Toolkit *interpreter*, which generates the corresponding costs according to the rules defined by a plug-in that has to be loaded when starting the system. The costs are then "consumed" by the accounting authorities, registered in the system as *context handlers*.

The plug-ins – pieces of software that are downloaded at runtime from the accounting authorities – define how the telemetry data are transformed into costs: they infer the road fee or calculate the accident risk and transform this risk to an insurance rate, for example. The raw GPS coordinates are transformed into meaningful location context using the location context auxiliary system. Every predefined period of time (a day for the prototype, a week or a month would probably be meaningful for a real-life deployment), the interpreter returns an aggregated sum to the context server, which in turn sends this sum to the accounting authority (over the vehicle's UMTS connectivity or the home WiFi network that can be received from the garage, for example). The fact that the accounting authority does only receive the accumulated sum but not the raw sensor data will be relevant for the privacy and security discussion in section 6.4.2.

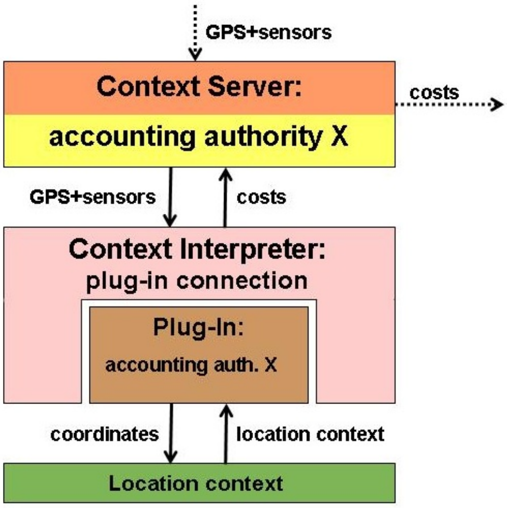

Figure 6.7: Plug-in-Architecture for the Accounting Authorities.

By using the Context Toolkit, adding new sensors or new accounting instances become easy tasks. To add a new sensor, a new *widget* has to be written that encapsulates the proprietary communication with that sensor. Similarly, a new accounting authority is being added by registering it as a new *handler*.

From a functional perspective, the Smart Tachograph knows three kinds of predefined accounting entities: `compulsory`, `required-select`, and `optional`. In a system configuration file, three corresponding lists have to be filled out. In a realistic setting, these could be for example:

```
compulsory:  vehicle-tax-authority; traffic-police
required-select:  liability-insurance
optional:  own-damage-claim
```

Every item in the “compulsory” list has to be active before the system can be started. An accounting entity is active when its server has been registered in the system as a subscriber for sensor values, and the plug-in has been downloaded from the corresponding authority server. Likewise, for every item in the “required-select” list, one plug-in has to be present. The difference is that the user may choose here between different plug-ins (e.g., from different insurance companies), while for compulsory plug-ins there is no choice. Any number of “optional” plug-ins can be loaded before starting the system, but none is required. In the prototype presented here, the Smart Tachograph software does not start until all mandatory servers have been started. In a real deployment, it is conceivable that the system would be connected to the electronic anti-theft device, so that the car would not start until all the legally required plug-ins are active. Also, in a real deployment, the tasks accomplished for the prototype by the laptop computer would probably be taken over by an on-board computing unit.

The software infrastructure accomplishes several other tasks as well. When starting the system, it presents the driver the accounting places that must be active to be able to start the car. These could for example be the tax authority and/or the police and at least one insurance company. For the first two there is no option, but the insurer may be freely chosen. A service description and discovery protocol is therefore part of the software infrastructure. Insurers may publish several insurance schemes they offer. These will typically have a pretty low minimum per-kilometer insurance price, which will increase according to traffic, time and place of the ride, momentary acceleration, driven speed in relation to the speed limit, and other criteria.

Driver Interface

The Smart Tachograph's software further includes a front-end interface to the driver (Fig. 6.8). The main system window, in the lower right corner, is needed to setup and start the prototype system. Several parameters can be set here. Among them, the driver may choose for all required but selectable items a specific instance. For example, he or she may choose a liability insurance from the existing offers. The driver may further choose any number of optional insurances.

To have an up-to-date view of the existing offers, the Smart Tachograph accomplishes the following two steps when the user starts the system (read: "enters the car" for a real deployment). First, it connects via the car's wide area communication system (e.g., UMTS) to a root-server, retrieving a list of available vehicle insurance companies. Then it connects to the server of each insurance company to retrieve the available insurance schemes, presenting them to the driver. After choosing everything needed, the user starts the Smart Tachograph by clicking the "start" button.

The second window (in the lower left corner) presents the sensor data as a collection of bars. The raw GPS coordinates are translated into the actual road that is been driven on using the location context subsystem. Knowing the speed limit for all streets and roads in the administrative region of Zurich, the system displays this information

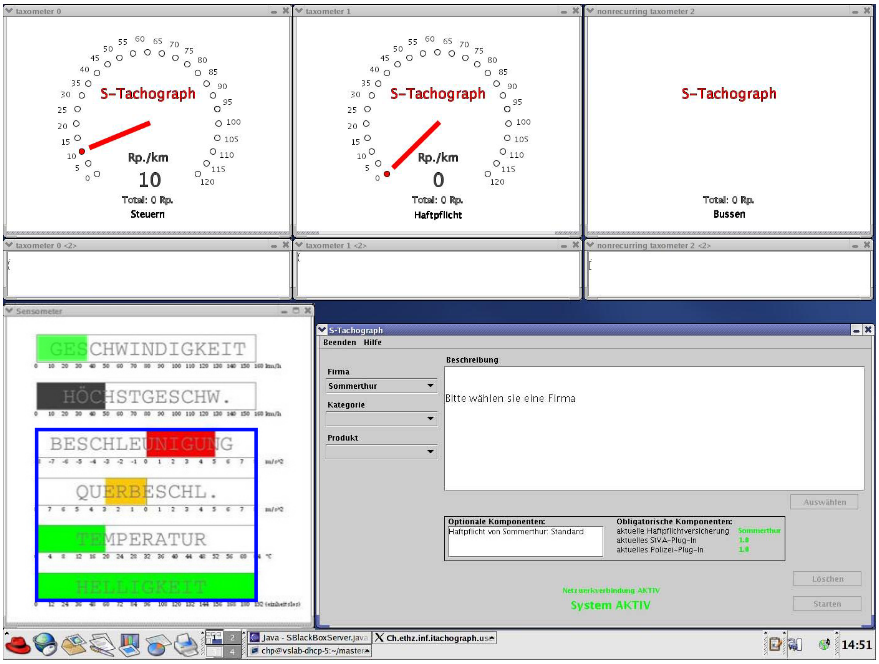

Firma: Sommerthur

Kategorie:

Produkt:

Beschreibung: Bitte wählen sie eine Firma

Optionale Komponenten: Haftpflicht von Sommerthur: Standard

Obligatorische Komponenten:
aktuelle Haftpflichtversicherung Sommerthur
aktuelles StVA-Plug-In 1.0
aktuelles Polizei-Plug-In 1.0

Netzwerkverbindung AKTIV

System AKTIV

Figure 6.8: A typical screenshot of the system's interface.

on the second bar from top. It uses this information, together with the actual speed (displayed on the topmost bar), for risk approximation. The lower bars show the data from the other sensors – longitudinal and cross acceleration, temperature, and light intensity. The data from all these sensors, as well as the GPS coordinates, are ascertained and transferred to the laptop computer once every second.

The third type of interface windows (upper part of Figure 6.8) would probably be the only ones shown in a real deployment while driving. They show the current costs (insurance rate, road tax), which are continuously calculated from the received sensor data. The indicators presenting these values should be perceived by the driver similar to the momentary gas consumption indication built into some cars. It allows the driver to receive instant feedback on how his or her driving habits influence the traffic costs. Traffic fines (window in the upper right corner) can also be issued automatically. They are not expressed as money per kilometer, but as one-time events (i.e., when the speed limit has been exceeded for more than ten seconds).

Context Toolkit

While designing the Smart Tachograph platform, we realized that we needed several components, most prominent a publish/subscribe mechanism for a flexible distribution of sensor events to the multiple accounting authorities. We also wanted means to interpret context according to certain rules. These rules had to be dynamically changeable to enable the system to cope with such cases as a change in the road pricing fees, or a frequent change of insurance companies, as suggested by [29] under the name of ride-by-ride insurance. We further wanted to be able to distribute the same result of the context interpretation to both the accounting authority and as real-time feedback to the driver. Finally, we wanted to archive the sensed data as evidence in case of a later dispute between the driver and one of the accounting authorities.

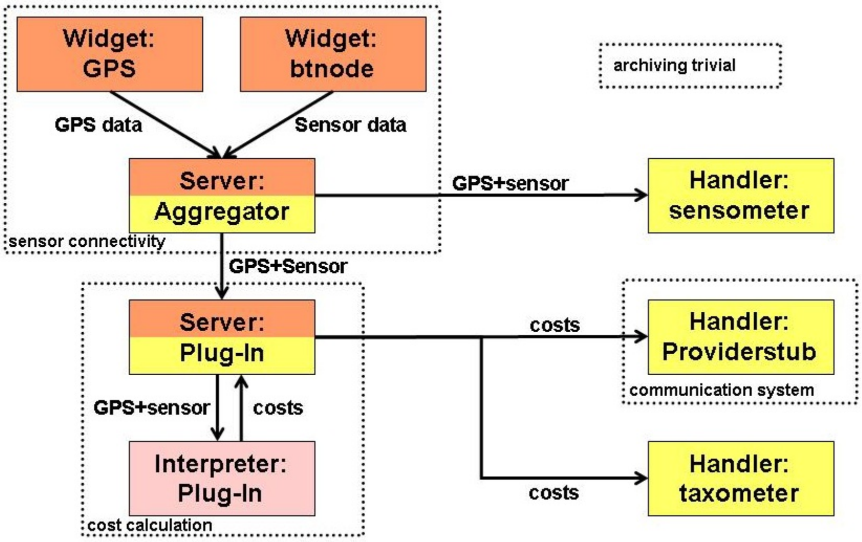

Figure 6.9: The Smart Tachograph's architecture in Context Toolkit terminology.

After looking at some available publish/subscribe mechanisms, we found Anind Dey's Context Toolkit [41] to be ideally suited for our needs, so the Smart Tachograph has been build on top of it. The Context Toolkit offers so-called "widgets"" to represent the context producers, "context servers" that receive context, transform it (usually generating higher level context) and pass the result further to other consumers, so-called "context interpreters" that modify context according to given rules and give the result back to the client, and "context handlers", that are context consumers. Moreover, archiving comes for free, since the Context Toolkit has by default a MySQL database archiving all context data. Figure 6.9 shows the architecture of the Smart Tachograph in Context Toolkit terminology. By using the Context Toolkit, adding new sensors or new accounting instances become easy tasks. To add a new sensor, a new "widget" has to be written, that encapsulates the proprietary communication with that sensor. Similarly, a new accounting authority is being added by registering it as new "handler."

6.3.4 Design – External Systems

As Fig. 6.2 has already shown, there are two sorts of external systems: the "root server," and the so-called "provider servers." Each of the latter is the service providing computing infrastructure of an individual accounting authority, the entity to which vehicles connect.

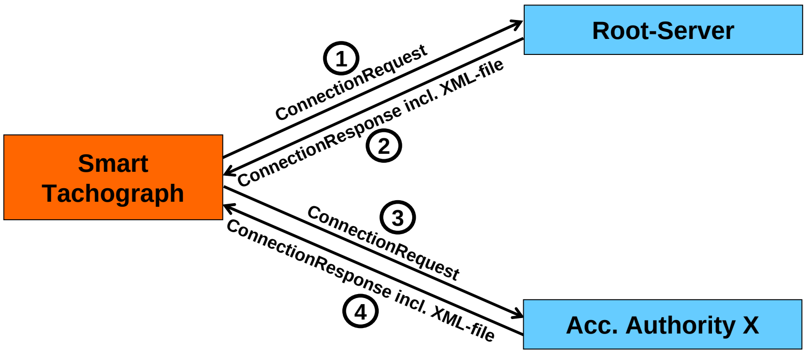

Figure 6.10: The first steps upon starting the Smart Tachograph: 1) It connects to the root-server, 2) receiving an up-to-date XML-file with the existing accounting authorities. 3) It then asks each of these for the current offers, 4) which are received as XML-files as well. Steps 3 and 4 will be repeated for every instance in the list originally received from the root-server.

Upon starting the Smart Tachograph system in a vehicle, it will first connect to the root-server. The root-server is the central authority (which, in a real deployment, would most likely belong to a government agency) keeping track of the existing accounting authorities. This is necessary since in a real deployment, as already argued in the last section, both the sorts of traffic costs billed through such a generic platform, and the set of authorities offering individual services, are likely to change. Step 1 in Fig. 6.10 represents this request to the root-server, which replies (in step 2) by a table of names and connection parameters (i.e., IP-address and port number) embedded in an XML-file.

Since for all networking purposes, accounting authorities are represented by provider servers, the next step for the Smart Tachograph will thus be to contact each provider server from the list (step 3 in Fig. 6.10) in order to receive their current offers (embedded in an XML-file as well,

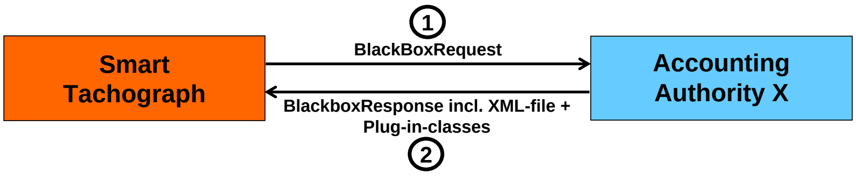

Figure 6.11: When the Smart Tachograph wishes to install the accounting plug-in from a specific authority, it requests it from the corresponding provider server. The plug-in will run on the vehicle's local computing environment.

step 4 in the Figure). For items in the `required-select` and `optional` lists, the driver will read the textual description of the offers and decide which one he or she wishes to install in the system. For every item he or she wishes to install, the communication as represented in Fig. 6.11 takes place. The Smart Tachograph requests a Java class from the corresponding accounting authority. This class is the accounting authority's local (i.e., on the vehicle's computer) plug-in for the Smart Tachograph's plug-in architecture as presented earlier, in Fig. 6.7. The plug-in accomplishes two tasks: First, it computes the traffic costs according to the rules that have been either imposed to the motorist, or that he or she agreed to when accepting the conditions and thus downloading the plug-in. Second, it sends these costs on a regular basis to its provider server.

As Fig. 6.12 shows, the system cannot be started as long as there have not yet been downloaded and activated into the system plug-ins for every accounting authority in the `compulsory` and `required-select` lists.[13] For `compulsory` sorts of traffic costs, the plug-ins get automatically downloaded by the system, for the `required-select` ones, the driver has to make the choices.

Finally, the system's sensors, although placed in the vehicle and logically part of the Smart Tachograph itself, are nevertheless external to the computer running the tachograph's software. Means to connect them to the software running on the vehicle computer had thus to be created. The resulting framework, called `btnode2java`, is sketched in Fig. 6.13.

For the BTnode computation and communication device, we have written a software called `carsensor_pull` (marked with an A in Fig. 6.13),

---

[13] The German text displayed at the bottom of the window in Fig. 6.13 reads: "System NOT ACTIVE, vehicle cannot be started."

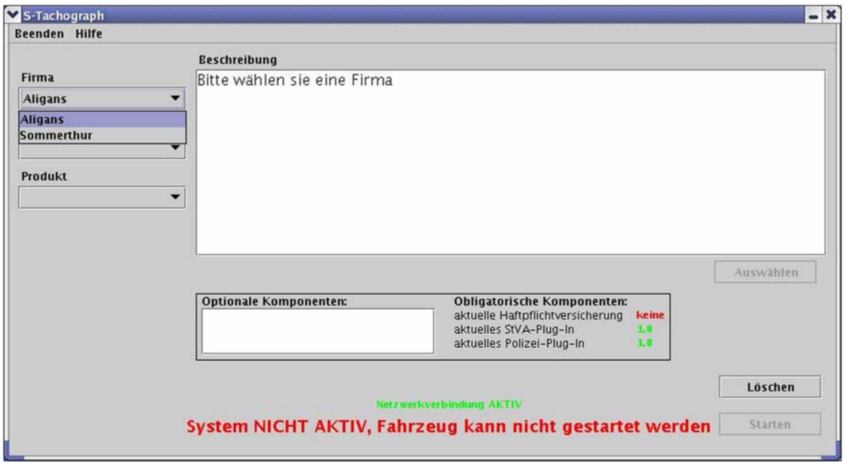

Firma: Aligans

Produkt:

Beschreibung: Bitte wählen sie eine Firma

Optionale Komponenten:

Obligatorische Komponenten:
aktuelle Haftpflichtversicherung keine
aktuelles StVA-Plug-In 1.0
aktuelles Polizei-Plug-In 1.0

Netzwerkverbindung AKTIV

System NICHT AKTIV, Fahrzeug kann nicht gestartet werden

Figure 6.12: Smart Tachograph initialization window. The `compulsory` plug-ins "police" and "road pricing" have automatically been downloaded (thus, they are highlighted in green). For the `required-select` item "liability insurance," the driver has not yet chosen an instantiation, hence the item is underlined in red in the lower right side of the window. The engine can thus not be started yet.

which can receive requests for sensor values via Bluetooth from the vehicle computer. `carsensor_pull` can run in one of two modes: continuous reading of the sensor data, and one-time, on-demand reading. In the continuous mode, if all sensors are polled, the system offers a frequency of around 1Hz, i.e., all sensor data are read once per second. At the other end, if only one specific sensor is being read, a frequency of about 7Hz is possible. The setting of which subset of sensors should be read in either continuous or on-demand mode is part of the API of `carsensor_pull`.

BTnodes typically communicate via Bluetooth among each other, but via a serial cable to computers. For a direct Bluetooth communication between the laptop computer running the vehicle's software and the BTnode connected to the sensors, a Linux-driver, called `btnode_sensor`, has been written (denoted B in Fig. 6.13). This computer-side driver sends the commands and receives the sensor data from the BTnode. Finally, to be able to integrate the so-gathered data into the Smart Tachograph, a Java wrapper-class, `BtNodeSensorAccess.java`, has been written.

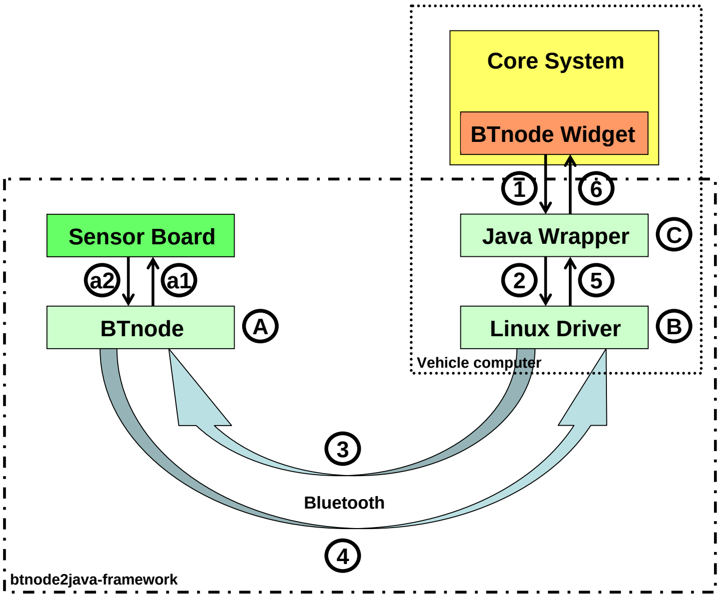

Figure 6.13: The `btnode2java` framework, connecting sensors to the Smart Tachograph system.

6.4 Analyzing Scenario and Prototype – Insights

Having developed the prototype, we set out to analyze possible implications of a real-life deployment of a system such as the Smart Tachograph. This corresponds to the third step of the scenario-driven prototyping method, as summarized in Fig. 5.2. Having both new scenario and demonstrator in our hands, we could analyze consequences on various technological and societal levels, as well as on levels that combine the two.

6.4.1 Practical Technological Limitations

During the development of our prototype, we have been out on the streets for several weeks, testing and tuning the system, gaining experiences about difficult driving situations that challenge the system, and partially solving those problems. We have also been an entire day on a closed circuit, where a professional driver tested the Smart Tachograph under different conditions – from "normal" driving to a driving style that would qualify as very aggressive and highly risky for an average driver [34].

The main problem we encountered was to ascertain on which street the vehicle is being driven on and what the speed limit on that street is. There have been several causes for this problem, first of all, the imprecision of GPS measurements. Especially within cities, due to the reflections caused by buildings, the GPS location determination can easily exhibit an error of up to around a dozen meters. Thus, often when looking up the position reported by the GPS sensor in the geospatial database<, the vehicle seems to be not on a street, but for example inside some building. This inaccuracy is easily solved by looking for a street not only at the precise position reported, but inside a circle around that position. However, the circle may not be too large, since several different streets could then be found, especially in the neighborhood of street crossings. A radius of eight to ten meters has proven to be a good trade-off. A street is always found and usually it is the correct one. For the few remaining inaccuracies, buffering the few last reported instances of the current street has proven to be a good workaround. A change of street is only accepted by the system if the new street is reported four times in a row by the geospatial database. Since the GPS coordinates are evaluated once every second, this buffering successfully removes annoyances induced by a short proximity to other streets, e.g., at crossings. It also helps to avoid a similar problem we observed when first testing the system. When driving on a highway under a bridge, sometimes the system suddenly changed from a speed limit of 120 km/h to 50 km/h. Since we were in fact driving at 120 km/h (the usual speed limit on Swiss highways), for one or two seconds the insurance price rocketed to a few dozen times more than it had been before. The system thought we were on the bridge and not the highway anymore, and this was the logical result of our "overspeeding" of 70 km/h. This kind of malfunction could also be successfully removed by introducing a delay before accepting a street change.

After ascertaining the street the vehicle is being driven on, the next point is to determine the speed limit for that street. This should be trivial in theory – a lookup in the geospatial database that contains, among other attributes, the speed limits for individual streets. In practice, however, this data is not easily available. After extensive research, we found only two providers of street topology data worldwide – Navteq and TeleAtlas.[14] All providers of GPS-based car navigation platforms seem to buy the raw geographical data from one of these two producers. Navteq has no speed limits recorded for Switzerland and the ones from TeleAtlas database were often incorrect. Although this is in part

---

[14] See www.navteq.com, and www.teletlas.com, respectively.

a rather pragmatic problem, which presumably occurs only for some geographical regions, there is also a conceptual issue behind it. Speed limits on individual streets change with quite a rapid pace, so that a CD containing them will be partly outdated from the first day of usage. Speed limits may change due to changes in the street architecture (e.g., a street enlargement may come with an increase in the speed limit), or traffic-policy reasons, but also because of short-termed construction sites. To have an outdated database in a system that could in consequence charge a sum one or two orders of magnitude larger than the true one, is obviously unacceptable. A prerequisite for such a system to work is thus to have an efficient way of propagating speed limit changes to vehicles. Many solutions are conceivable, but they all require quite a massive infrastructural support, which is unavailable today. A centralistic solution could be for example to have a publicly-accessible database where all changes are published. Vehicles would lookup that database on a regular basis, updating their local copy. A distributed way would be to have electronic tags on all speed limiting signs, which could be read by the vehicles.

Looking at realistic applications and only estimating the risks and opportunities of technology that is likely to shape our future, is one of the primordial aims of scenario-driven prototyping. Pragmatical issues of this kind can only be discovered by building a prototype that is as close to reality as possible and testing it in real-life settings. Thus, discovering and weighting the relevance of such technological limitations, albeit not of particular theoretical interest, are nevertheless important steps for estimating the degree of realism of a certain application, and subsequently for the spirit of scenario-driven prototyping.

6.4.2 Privacy Considerations & Client-Side Personalization[15]

When conceiving a system such as the Smart Tachograph, the most obvious societal concern is the high degree of possible intrusion into the driver's private sphere. The straightforward technical approach, having all driving data sent to the accounting authorities, would reveal to insurers, tax authorities, police and others an unprecedented amount of information – precise data about the whereabouts, speed and accelerations of the vehicle at every moment, even when it is not moving. For the family car, possibly not only these data would be taken

---

[15] Parts of this section are based on [28] and on joint work with Marc Langheinrich [35].

into account, but also which member of the family has been driving. Such data would not only be useful to insurers or the tax authority, but also to marketers, neighbors, political enemies, or law enforcement. Early within the development of the Smart Tachograph the question thus raised: Can information asymmetry within the vehicle insurance market be broken to the benefit of safe drivers, insurers, and overall economic output without having to give up virtually all privacy?

After a thorough analysis, we realized that the architecture of a system that continuously analyzes the driving parameters to ascertain momentary costs on a pay-per-use basis can be realized in three different ways. The first, chosen by some insurance companies and which is also the standard for road pricing schemes (see related work in Section 6.5), is the above-mentioned straightforward technological solution to send all sensed data to the accounting entity, be it tax authority or insurance company. The data could be sent online (via the car's wide area communication system) or offline, on a regular basis. This solution is the most simple yet most privacy invasive.

A second possible implementation – at least for the insurance part – would be to not disclose the data by default, but store it and reveal it in order to get a retroactive reduction for a safe driving style. The data could be stored either locally in the car's blackbox, or it could be encrypted with the motorist's private key and transferred to the insurance company. To qualify for the reduction, the customer has to reveal the data to the insurer. This is the model chosen by the Progressive insurance company for one of its products, as section 6.5 will show. Aside from being a different insurance model than the risk-dependent momentary insurance rate presented here, this solution has several drawbacks. As pointed out in section 6.5, in order to get the price advantage, the customer has to reveal all data and thus give up privacy in this model, too. Furthermore, the responsibility lies with the customer, and it is unclear what the long-term consequences of both revealing or not revealing the data would be. What will for example happen when the contract is up for renewal and the driver has chosen not to send "enough" data to the insurer? Finally, it is unclear how this approach could work for a road pricing scheme, being thus rather unsuited for a generic solution for all types of traffic costs.

Client-Side Personalization

In the third and chosen model, we tried to provide evidence that highly personalized insurance rates are also feasible without such a significant loss of privacy. Instead of sending the telemetric data to the individual accounting authorities, all data is processed locally and only a gross total is transferred periodically (i.e., once a day, or once a month), as described in section 6.3.3. The data is processed by a plug-in that has been downloaded from the respective authority (technically a Java-class). Such a *client-side personalization*[16] insurance scheme [35] guarantees a high level of privacy. By sending only the gross total at the end of the billing period to the accounting authority, not only do the whereabouts of the vehicle during this time remain undisclosed, no inferences on the individual driving style can be drawn either. Even if a driver ends up with a high rate at the end of a period, this sum may be the result of a risky driving style, but also be caused by a very cautious, yet heavy-mileage driver. Past whereabouts and behavior of users are protected, yet they pay their fair share. Moreover, since the sensor data needs not to be cached (the continuously incremented overall sum will suffice) or will only be cached on client-side, the driver retains full control over his or her data, which reduces the possibility that data gained by sensors will be used against the driver.[17]

However, as we have analyzed in [35], a client-side personalization scheme poses a range of new challenges. They ultimately require the system to be designed in such a way that neither party may tamper with the system (or at least that such an attempt will be noticed). Customers need to be assured that the downloadable code bills them according to the agreed-upon contract, while insurers must be satisfied that the system will not allow customers to submit false (i.e., lower) totals.

System Challenges – Insurers Safety

Drivers could attack the system in several ways. An obvious one would be to modify the software class S received from the insurer into a new software piece S', which then sends an insurance rate R' lower

---

[16] Client-side personalization is used in the area of user modeling and systems personalization to respect privacy despite personalization, e.g., [21, 118]. We have proposed to apply the concept within ubiquitous computing, specifically for the individual accounting of traffic costs.

[17] Of course, this does not preclude court-orders that might force the driver to turn such data over to law enforcement, in case the data has been cached on client-side.

than the original rate R that would have been sent by S. The system thus needs a mechanism for tamper-proof software distribution, which in turn is one of the core functionalities of the Trusted Computing Platform (TCP).[18] Originating from early work on secure bootstrapping [8], TCP-based systems will allow software distributors to securely deploy software programs by using special hardware chips to secure a machine's initial state, memory, and computation, as well as having specific software modules in both the operating system and the application software. While the deployment of trusted computing is highly controversial in the area of consumer devices (e.g., personal computers, media players) [149], such an infrastructure can play an important part in safety-critical systems, where there is less of a threat for curtailing individual rights such as fair use. For our purposes, the implementation of a TCP inside the black box will effectively ensure the correct deployment of the insurer's billing code.

Given a correctly functioning software S, however, the vehicle's owner may alternatively try a man-in-the-middle attack by intercepting the message C from S to the insurer and replacing it with another message C'. To counter this, we require a black-box-specific secret key s_c, which then allows us to digitally sign all messages sent by S. Such a key might either be part of the TCP on the black box, or a customer-specific key given out by the insurer (e.g., as part of the software distribution step). Given the sealed storage property of TCP, this key is not accessible from outside our black box, thus making it impossible for the vehicle owner to replace the message C. By including a timestamp or serial number with each message C we can also prevent replay attacks (i.e., replacing a message C_h reporting a high rate with an older message C_l that reports a lower rate).

Last but not least, a driver could try to tamper directly with the vehicle's sensors, e.g., inserting a device that would cap the reported speeds to always stay below a certain (safe) limit, or simply covering up the rain sensor, so that it never reports rainfall (as driving under rainy conditions would make the insurance more expensive). Note that this is not an attack specific to the client-side personalization paradigm; it would be equally effective in a system reporting all sensed data to the accounting authority. However, having all data reported and employing data mining algorithms, similar to those used by credit card companies to empirically discover a high probability fraud, insurers might be able

---

[18] See https://www.trustedcomputinggroup.org/home.

to discover such tampering based on extensive analysis of valid driving records. For example, a security system on the insurer's side could issue an alarm if a certain change speed would be accompanied by certain non-standard acceleration changes, or if the reported telemetric data does not follow the typical values at the car's current position.

To counter a sensor tampering attack, we see a number of possibilities. The most obvious and probably least realistic would be to include such detection heuristics into either the operating system of the black box, or alternatively into each insurer's billing software. Unless such detection methods could be boiled down to some simple algorithms, the limited computing and storage facilities of the box might render this infeasible. A second, less powerful alternative would be to have the box periodically send position-neutral telemetric data to the insurer, in order to at least cross-verify the various sensors for acceleration, speed, temperature, and road type. The rogue customer taping over the car's rain sensor might thus get caught as the reported dryness would not match the also-measured humidity levels or the traction feedback from the wheels. More promising might be the approach of using data perturbation techniques from the field of statistical databases [1] to randomize the driving record before sending it off to the insurer, thus allowing for a statistical validation of the telemetric data without disclosing individual trips. One might further envision building equally tamper-proof sensors that would communicate with the black box using encrypted and authenticated channels, though it might still be several years before this would be cheap enough. An even simpler solution may be to include as many sensors as possible into the tamper-proof black box hardware, although this approach may not work for some of the sensors (e.g., the rain sensor).

System Challenges – Customer Safety

Equally imperative for the acceptance of such a system will be the customer's ability to verify the correct functioning of the billing mechanism, i.e., that the daily or monthly charges correspond to the agreed-upon insurance contract.

Rogue insurers might be tempted to simply charge a sum C' different (i.e., higher) from the 'true' sum C reported by the car's box. Note that while market pressures might prompt most insurers not to cheat in such a blunt way, the ability to quickly switch insurers with such an infrastructure (potentially on a trip-by-trip basis [29]) might prompt developments similar to today's telephony market, where many consumers simply choose the cheapest provider on a call-by-call basis, thus risking to fall prey to scrupulous fraudsters. To this extend, customers might want to verify not only the values that the unit is reporting back to the insurer, but also keep a local log of these sums in order to verify the totals billed. In order to allow verification of the encrypted total C, our previously stipulated secret key s_c must be part of a public key cryptography system. This enables a customer to decrypt this value using the corresponding public key p_c (which would be made available together with the installed black box or the downloaded billing module). Equally important would be to have the billing software also available in unencrypted form in order to allow customers to locally verify the computed values.

The second customer-safety requirement concerns the correspondence of the downloaded software to the advertised rates. One option would be to standardize the computation in such a way that insurers would only publish and download the specific parameters. However, this might not only restrict the insurers' ability to differentiate their products sufficiently, but also turn out to be too simple to properly compute the correct risk value of a specific driving style. Alternatively, consumer interest groups or government bodies might offer a testbed-environment in which published billing algorithms could be tested under a variety of simulated conditions (potentially using tools from formal software verification to fully automate this task) and which would award trust-seals to verified algorithms. Note that such open-source billing algorithms would not necessarily infringe on an insurer's trade-secrets, as these typically lie with the way these algorithms or parameters are chosen, rather than with the published rates themselves. Also, as will be addressed shortly, our research has shown that many insurers would prefer rather simple algorithms based on few parameters only, in order to better communicate (i.e., sell) these to their customers.

6.4.3 Computing the Risk

Obviously – even when having a multitude of fine-granular, real-time parameters such as cross and longitudinal accelerations, driven speed related to maximum allowed speed and type of street, the time of the day, weather conditions, traffic, and so on – assessing the exact risk from such parameters is a non-trivial task. Our prototypical insurance plug-in of the Smart Tachograph takes into account five parameters: type of street, difference between driven speed and speed limit, the two acceleration types (longitudinal and cross acceleration), and the time of day. For converting these sensor data into an accident probability and expressing the risk in a monetary way, we made the following assumptions: There is a basic per-kilometer risk, that depends on the type of street and on the time of day, as suggested in [119]. The per-kilometer accident risk is lowest on highways, followed by country roads, and is highest in cities. It varies between 2 and 10 cents per kilometer. In a real deployment, this minimum would presumably also depend on the "classical" risk factors such as driver age or experience, which are used today to classify customers into driver categories.

With respect to acceleration, we acknowledge that some thresholds have to exist. At every traffic light stop, every departure, and every curve taken by the driver, there are accelerations involved. Such accelerations within normal limits pose no special danger and have to be allowed without punishment. According to our subjective danger sensation correlated to the measured accelerations, the thresholds for the prototypical plug-in have been set to $2m/s^2$ ($1/5g$) for the cross accelerations as well as the positive longitudinal acceleration, and to $3.5m/s^2$ for negative longitudinal acceleration (braking). After that, we assume that the risk increases exponentially with a low base of 1.5. We further assume that exceeding the speed limit with 10% does not notably increase the risk, and only after that it increases – exponentially as well, but with a lower base than in the case of acceleration. This base further depends on the street type, varying between 1.05 for highways and 1.2 within cities. The overall formula for the momentary insurance rate for the experimental system (expressed in cents per kilometer) thus results in

$$R = B_{st,t} * 1.5^{Max(0,\frac{A_c - 2}{2})} * 1.5^{Max(0,\frac{Abs(A_{ln}) - 3.5}{3.5})} * VC_{st}^{Max(0,\frac{V_{driven}}{1.1 * V_{limit}} - 1)} \tag{6.1}$$

where R is the resulting momentary insurance rate, B the basic rate depending on street type and time, A_c the cross acceleration, A_{ln} the negative longitudinal acceleration (actually, this is a slightly simplified version of the formula, ignoring the positive acceleration), VC the base for the speeding coefficient, V_{driven} and V_{limit} the driven speed and the speed limit, respectively.

Subsequently, we had two rounds of discussions with stakeholders from the insurance industry. The first talk, on May 3rd, 2005, has been with Mr. Anton Brunner and Mrs. Bettina Sinzig, from the former Winterthur insurance company.[19] Mr. Brunner is head of the accident research group, Mrs. Sinzig being one of his colleagues. We have learned that since the early 1980s, large insurance companies such as Winterthur have established a strong accident research tradition. They have been doing countless measurements over many years and know with quite a high precision when driving style parameters, such as the involved accelerations, become dangerous. Up to now, they have not been able to transfer this know-how to pay-per-risk insurance schemes due to the bulky equipment and the high costs involved.

During the discussions, it became clear that some of our assumptions have been quite exact, while other have been rather erroneous. The experts confirmed that a defensive driving style usually remains under $0.3g$ of cross acceleration, grows exponentially, and becomes really dangerous around $1g$. At $1g$ cross acceleration, our formula results in an insurance fee almost 8 times higher than the basic fee (and then continues to grow exponentially). Choosing an exponential function seems to have been the right decision, details aside. However, for the braking we have set the threshold clearly too low. According to Mr. Brunner, below $0.8g$ of negative acceleration, quite a normal braking process takes place. Braking above that threshold, however, is a clear indication for an emergency braking. The threshold has thus to be higher than we set it, but then the curve of risk has to grow faster than it did in our original prototype.

On June 9th, 2005, we had a further round of discussions with representatives from another major Swiss vehicle insurer, the Zurich company.[20] Mr. Jürg Kufer, head of the "development of new customer services" division, confirmed that with most assumptions of the insurance formula, as well as with the fact that insurers would have a high interest to offer such individualized insurances, if technologically possible and economically feasible (i.e., low-cost), we were on the right track. Our formula seems to have over-estimated the velocity component though. Overspeeding is one determinant for accidents, but this depends much more on the actual context than, for example, with accelerations. Often, speed limits are set very conservative or with respect to other criteria (such as noise reduction), thus overspeeding may

---

[19] Today Axa-Winterthur, see www.axa-winterthur.ch.

[20] www.zurich.ch.

often have no influence whatsoever on the risk of causing an accident. Since the correlation between maximum allowed speed and the accident danger in case of exceeding that speed is very volatile we came up with a pragmatic compromise. We have increased the allowance from 10 to 20% of, but to at least $20km/h$ over the speed limit. Below this new threshold, there is no punishment, above it the insurance costs start to slowly grow.

However, during the discussion it became clear that there in no generally correct way of expressing the influence of overspeeding on the risk, since this influence depends on too many further contextual factors. This plainly shows some conceptual limitations of our system. Although it is seemingly just a detail, with no strong overall influence on the accident risk, it demonstrates that reality is sometimes inherently so complex that it cannot be quantified into computer systems. This observation will probably be relevant for developers of other ubiquitous computing systems as well: No matter how advanced and varied the sensors automatically measure factors of the physical world (such as accelerations and speed in our prototype), and how smart and complex the programs transform this raw into meaningful data (accelerations into likelihood of accidents, and thus money), reality will always remain too complex to be completely and correctly represented in computer systems.

Finally, another parameter to be considered, which according to the industry representatives would clearly be more relevant than driven speed in correlation to speed limit, would be the distance to the car in front correlated to the driven speed. Since we had no immediate means to measure such data, we could not include this in our prototype.

6.4.4 Acceptance

After taking all these steps, we not only had a low-cost, easily scalable solution for breaking information asymmetry and thus technologically enabling individual, risk-dependent insurance premiums, we also had now a pretty clear imagine of how the formula for calculating these risks had to look like. We had further been able to discover and largely get rid of the pragmatical technological problems presented above, and – more importantly – to come up with a solution that – profoundly different than all other insurance prototypes that will be presented in the next section of related work – succeeds at respecting the customer's privacy while at the same time charging him or her the correct costs.

However, in the discussion with Mr. Kufer from the Zurich insurance company, a new acceptance issue has been pointed out, which leads to an inherent trade-off between acceptance and exactness. Requesting customers to sign a complex mathematic equation such as Formula 6.1 in the previous section as part of an insurance contract seems to be unacceptable, no matter how much a realistic risk estimation it encompasses. Insurance companies know that a "good" contract (i.e., one that can be accepted by customers) must comprise two features: simplicity and upper bounds. To keep complexity low, according to Mr. Kufer at most three real-time attributes should be considered. Because of their outstanding importance, these could be cross acceleration, the distance to the car in front, and the number and intensity of emergency brakes. To further reduce the complexity perceived by the customer, discrete intervals are preferable over a continuous function. There could be, for example, three different classes of driving: safe, normal, and dangerous driving. After a journey, the motorist would get different per-kilometer prices for the times he or she exhibited any of these styles. This, however, leads to the interesting situation where an insurer would be technologically able to offer an insurance premium closer to reality, and for a large group of customers it would be profitable to accept it. However, due to the sheer complexity and the resulting insecurity on customer side, they would not accept such a contract, so the insurer is forced to limit the precision and amount of its measurements.

Not having defined upper bounds in the insurance tariff, on the other hand, could lead in extreme circumstances (when the sensors measure that an accident is bound to happen) to a chronic situation in which the insurance rate would get as high as the expected damage costs. This would, of course, undermine the idea of an insurance, making such a practice unacceptable. But even in less extreme situations, customers seem to feel much more at ease if they know what to expect in terms of maximum possible costs, no matter of their behavior.

Combining these acceptance measures – only measuring the few most relevant parameters, having discrete risk intervals instead of a continuous function, and imposing upper bounds – represents, however, a step back from the outmost precision that has been at the very core of the Smart Tachograph vision and project. And, more importantly, it shows another possible limitation in the deployment of ubiquitous computing technology in general. Not only for privacy reasons and because of the fear of being treated unjust (as in the example of the amazon bookstore selling DVDs for different prices) would customers reject a new application. Complexity can also be a reason for rejection. What stands to reason for obviously noticeable complexity, seems to be true for hidden complexity as well. In the Smart Tachograph example, the complex risk estimation formula would be hidden from the customers, yet they would feel at unease nevertheless, because of not being able to understand how the resulting sum came to being. Since, however, ubiquitous computing's very aim is to add complexity (even if in a way hidden from people's eyes), this seem to possibly lead to some acceptance issues not anticipated by Mark Weiser when he wrote his optimistic vision of the future of ubiquitous computing and the way it will free people from tedious tasks [163].

6.4.5 Further Economic Issues

We have argued above why not imposing upper bounds to an insurance contract could lead – precisely because of the precise measurements enabled by it – to situations that would contradict the very idea of an insurance. A common misconception, on the other hand, is that highly personalized insurance rates by themselves, which finally lead to "risk communities of one" (as Andreas Schraft, Head Risk Engineer of SwissRe, a reinsurance company, put it in a recent talk) would undermine the idea of insurances as well. This often-heard interpretation originates in the (wrong) assumption that an insurance company has to insure the same risk for a large number of people in order to work. As a matter of fact, insurance companies insure unlikely, but high-cost events of individuals with a sum that represents the probability of that event's occurrence times the costs it will produce (plus the insurance's security margin and its profit). Since many individuals are insured with one company, some of these events will occur, most will not, and due to the law of large numbers the company will with large probability pay the expected cumulated costs (if it estimated correctly the individual risks). While the company does need a large number of individuals for the system to work, it does not need to insure the same risk for everyone – highly personalized risk insurances will do just fine.

Even if, due to acceptance issues, individual pay-per-risk insurance schemes still lie a few years in the future, there are some areas more likely than others to be early adopters of such technology on a large scale. All domains where people drive cars that do not belong to them, but are merely leased, rented, or co-owned, fall into this category. In such areas, the driver of the car has an explicit or implicit obligation to handle the confided vehicle with care. And the lending or owning part can more easily enforce a black box analyzing the way customers handle the assets. For example, a US car rental company has started as early as 2001 to protect its vehicles from overspeeding by charging customers exceeding 70mph with a high fine.[21] Car sharing models, such as the popular Mobility-network in Switzerland,[22] have to pay relatively high insurance rates because of the few accident-prone risky drivers. Eventually (after two or three accidents in a short period of time) these drivers are sorted out, but the damage is done, and through the continuous flow of new members the insurance rates stay high. Car sharing networks would presumably be happy to detect such risky drivers before they cause accidents, by having their driving style analyzed from the very first ride. And the "How am I driving?" sign could soon disappear as well, if companies would start equipping their car fleets with black boxes analyzing (and reporting) the way their drivers behave. Furthermore, in such areas where different drivers use the vehicles only for short periods of time, the driver trying to tamper with the system seems less probable for two reasons: the expected gain is smaller (because of the short usage time); and the car returns periodically to its owner, who may detect the fraud. Thus, a lower level of security measures could suffice, further increasing the chance of early adoption.

6.5 Related Work

Within the ubiquitous and pervasive computing community, there has been little research regarding road pricing or pay-per-risk insurances. There have been, of course, numerous location-awareness projects in the broader domain of traffic, most prominent CoolTown's WebBus [88], but none that we are aware of regarding specifically road pricing or dynamic insurances. From inside these communities, our work has mostly benefitted from context awareness research, the Smart Tachograph platform being built on top of Anind Dey's Context Toolkit [41].

Some economists and business analysts [50, 119], however, aware of

---

[21] See http://archives.cnn.com/2001/TECH/ptech/06/22/gps.airiq/.

[22] See www.mobility.ch.

the potential of ubiquitous computing technology, have examined the impact of ubiquitous and pervasive computing on insurance markets from an analytical point of view, specifically highlighting the vehicle insurance market. All the rest of the related work comes from outside computer science as well – it is to be found in economic and business publications on road pricing or vehicle insurances.

A good overview on road pricing systems in use today can be found in [47]: the London Congestion Charge, Singapore's ERPS, the systems in Oslo and Trondheim, and California's expressway SR91. All systems but the one in Singapore have a very coarse area model. They either penalize the usage of one road only (the California expressway), or the entrance to a specific area, typically the city center (London, Trondheim, Oslo). In all these examples, the fee is flat for a multi-hour usage permit, typically for 24 hours. Obviously, this only allows a coarse traffic management, keeping some traffic out of the surcharge area, but having much of that traffic redistributed on the borders of the restricted area. Moreover, such a flat system does not acquire a high degree of fairness, since everybody pays the same independent of the actual usage. In Singapore, the model is more complex, every usage being charged and the fees varying from street to street and with the hour of driving. The more jam-prone the street and the less fluent the traffic rolls, the higher the tax gets. Singapore's road pricing has a very fine-granular model, however, it can not include other parameters, such as ozone or carbon dioxide levels. Another drawback of all these systems, as compared to the Smart Tachograph architecture, is that they rely on a heavy infrastructure of active RFID tags, ubiquitous gates with powerful antennas to read those tags, numerous cameras to identify the cheaters, and partially also payment machines and manned stations – they are thus substantially more reliant on heavy infrastructure, while at the same time considerably less nimble and extensible.

Early discussions in the field of pay-per-risk insurances revolved around considering driven mileage only. Litman [102] made the case for including the driven mileage into the calculation of the insurance rate, also pointing out the positive environmental and traffic safety side-effects, but ignored other criteria. Progressive, a US-insurer, initiated a pilot project (called Autograph) between 1998 and 2000. It took into account the driven distance, and in addition the time of day and the geographic location.[23] More recently, Norwich Union, a UK-based insurer, offers

---

[23] See www.epa.gov/projectxl/progressive/index.htm.

a black box for what they call "pay-as-you-drive" insurance,[24] but disclose only that they take into account the hours of driving, not whether their black box also considers other attributes. Privacy does not seem to be a central issue: "The black box device measures vehicle usage and sends data directly to Norwich Union using similar technology to that used by mobile phones." Progressive also started in 2006 to offer a more sophisticated insurance product, called TripSense.[25] TripSense is based on a black box that has to be installed in the car as well, but their webpage is more detailed about what it will record: "[. . .] which measures your actual driving habits and allows you to earn discounts on your insurance by showing us how much, how fast and what times of day you drive." The driver may analyze the data recorded over several months at his or her PC at home (and see, for example, the per-day number of "aggressive brakes" and of times driving over 75mph) and decide for himself or herself on sending the data to the insurance company or not. If the data is not sent, a no-punishment policy is advertised. This seems to be a more privacy- and customer-friendly approach than Norwich Union's, although in order to gain the price advantage, the customer has to send all data and thus give up privacy here, too.

The more important point, however, is that the responsibility lies with the customer. He or she has to decide on sending the data to the insurer, without possibly realizing what the longer-term consequences of such action will be. What will happen at the next contract renewal if the driver has not sent any data to the insurer in the previous period of time? What if the data was sent and points to a risky driving style? Who else will gain access to the data and will "my" data ever be used against me? Progressive states that: "We may retain the information that you send to us indefinitely" and further that "If you are in an accident, you may have a legal obligation to preserve the information on the TripSensor. This information may be sought by opposing parties in a civil lawsuit or by police when investigating the cause of an accident. We may be legally obligated to provide such information in response to a subpoena or as otherwise required by law."

To sum up, neither of the above-mentioned systems have been built with customer privacy at its core. All presented road pricing infrastructures record the places where the vehicle has been (or at least where it entered and exited the fee area) through transponders, while the exist-

---

[24] See www.norwichunion.com/pay-as-you-drive/.

[25] See https://tripsense.progressive.com/home.aspx.

ing insurance models continuously gather data about the driver's habits and whereabouts and send them to the insurance company. With the Smart Tachograph, we try to provide evidence that highly personalized insurance rates are also feasible without such a massive loss of privacy and control.

Also, we know of no approaches so far that tried to develop an open platform that could be used to calculate and charge a great variety of traffic-related costs. Proprietary black boxes have been the standard solution so far for pay-per-risk insurance prototypes; transponders and a heavy gate infrastructure have been the standard for road pricing schemes; but the two concepts have not been combined before. Summarizing, we believe that the power of our approach lies in its simpleness, flexibility, and open character.

6.6 Summary

In this chapter we have presented a first project in which we have applied the paradigm of scenario-driven prototyping introduced earlier, in chapter 5. In a first step, within an interdisciplinary dialogue, we have identified an area that could quite possibly be entirely reshaped by the deployment of ubiquitous computing – the propagation of individual and dynamic prices into various economic areas. The technological means exist already or are being developed; the business interest to introduce a new pricing paradigm which is customer-, context-, and behavior-dependent exists as well.

In a first future ubiquitous computing scenario, we have presented this paradigm in the context of a supermarket where the products continuously adjust their prices to prevent depreciation and to maximize turnover. In the subsequent discussions with economists and technology philosophers, this application domain has been criticized as being unrealistic. Customers would most likely not accept such economically and technologically unpredictable supermarket shopping, where past experiences would be worthless. The feeling of unfairness, namely that others might pay less for the same products due to some obscure reasons, would only add to the lack of acceptance. Another reason making such application domain unrealistic is the overly complicated and expensive infrastructure (as compared to the relatively cheap supermarket items) that would be needed.

The core idea of scenario-driven prototyping is to focus the technology assessment efforts exclusively towards domains and applications which to the best possible knowledge today are technologically, economically, and societally likely to emerge. Accordingly, we have changed the scenario almost in its entirety, keeping nevertheless the core idea of dynamic and individual pricing. After thoughts on which type of goods are most likely to be forerunners of the new pricing paradigm, in further discussions with economists, we have spotted a novel domain noticeably more likely to soon offer highly individualized prices – traffic costs, and especially vehicle insurances.

The new scenario envisions a system for the precise, individual, behavior and context dependent accounting of all sorts of traffic costs, many of which nowadays are spread across society or evenly (and thus, unfairly) shared between a larger group of motorists. Since charging individual motorists different amounts of money would not be discretionary, as in the original scenario, but related to their actual road usage and driving style, and since vehicle insurances already spread drivers into classes of risk, paying different premiums, the system's acceptance seems a large deal more likely than with price-changing supermarket items. Moreover, as a side-effect, the large-scale deployment of such a system would also have positive traffic safety and environmental effects.

Further following the paradigm of scenario-driven prototyping, we have developed a prototype, the Smart Tachograph, whose design has been thoroughly presented. The Smart Tachograph tries to be as close to a real implementation as possible. No aspects are simulated; we have deployed real sensors, gathering in real-time data about the vehicle's position, the maximum allowed speed at that specific point, the vehicle's longitudinal- and cross-accelerations, and a few other, less relevant parameters. Furthermore, not only the vehicle-internal aspects of the Smart Tachograph have been developed, the external systems – mainly the entities accounting traffic costs, called accounting authorities – have been developed as well. However, the main features of the Smart Tachograph are twosome. For one, it is the system's flexibility and general scope: New sorts of traffic costs, new sensors relevant for computations, new formulas for computing already existing types of costs can all be added with ease. The system can thus be used for any type of traffic costs, providing solutions for at least the two discussions that have been pursued increasingly, but always disjunct, in the last years: individualized insurance premiums, and fair road pricing according to the originator principle. The second feature is the used paradigm of client-side personalization. Using it, we have been able to prove, as opposed to all other behavior-dependent prototypes developed in parallel by the insurances industry, that a highly individual accounting of traffic costs is possible with a minimal level of privacy intrusion.

Nevertheless, the central point for being able to make a statement about the quality of scenario-driven prototyping as a tool for the assessment of ubiquitous computing implications, is the quantity and quality of the insights gathered along the process. A first indication has already been the change of scenario from a probably unrealistic to a substantially better suited domain. Without letting the expertise of social scientists flow in at an early point of development, we might have fallen into the pitfall many ubiquitous computing projects stumble into. We might have developed an interesting gadget, nicely to demonstrate, but actually irrelevant, and even dangerous if misleading public, politicians, and journalists to a wrong, unrealistic future path for ubiquitous computing.

Towards the end of the chapter, we have presented several other insights gained from the project, which, taken together, emphasize our claim from chapter 5. When presenting our approach, we have claimed that scenario-driven prototyping not only unifies the advantages of the two approaches used so far for the assessment of ubiquitous computing implications, but also allows for new insights to emerge from the synergy of the two. Technological insights have shown, for example, that maps with information as precise as needed for such a system are not available nowadays and would need years to be developed. And they would still not be able to cope with the dynamics of street changes. Other solutions would be required, all needing quite some amount of infrastructure.

Further stakeholder discussions have revealed some surprising insights, most of which we believe are generally valid for any ubiquitous computing application. Although being able to measure real-life events with unprecedented precision, for example, some parts of reality seem to be so complex and their interpretation so context-dependent that they can inherently not be grasped within a ubiquitous computing model of reality. Too much information, on the other hand, even if possible to gather and interpret, seems to lead to lack of acceptance from customer side. Not necessarily because of privacy concerns (which have been largely eliminated in our prototype), but simply because the sheer amount of data and the cryptic algorithms used by the system make people feel insecure.

This last example could not have emerged without both a prototype that showed how much information can really be gathered, and the cross-disciplinary analysis including stakeholders. It stays thus as proof for both our claim that some of the insights could not have emerged without combining engineers' and social scientists' know-how, and for how societal insights may influence technological developments, and technological insights may in return make the societal analysis more focused on relevant aspects, as discussed in the "reality-check cycle" of Fig 5.1. For the above-mentioned example, the analysis revealed that too much information and too precise formulas would be unacceptable for customers. They would, however, be needed for a measurement as precise and as individual as possible, which would in turn represent an advantage to a large number of customers.

Revealing such seeming contradictions and trade-offs at an early point into the development of ubiquitous computing applications, might make the whole difference in perception: between Weiser's vision of "calm technology," and the opaque, uncontrollable, feared "work of the devil."

7 The Chatty Environment – Enhancing the Everyday Life of Visually Impaired

As opposed to the technology-driven Smart Tachograph scenario, the Chatty Environment evolved inside the Ladenburg collegium as a typical problem-oriented scenario. Understanding from the early stages of the project the large potential of ubiquitous computing in the area of assistive technologies, and being able to get first insights into the everyday problems of blind and visually impaired from a blind colleague, we decided to develop a scenario aiming at solving or at least alleviating some of these problems. During our first investigations, it quickly became clear that blind and visually impaired encounter numerous barriers in their everyday lives, many more than the sighted researcher could possibly imagine. While some of these problems have partially been solved by already existing technology or legislation, the vast majority still remains unaddressed. At the core of the "independent living" scenario[1] thus lies the question of how the upcoming ubiquitous computing technologies could be used to alleviate some of the problems encountered by the blind.

As with the other project, here too we followed the scenario-driven prototyping methodology, in a process that heavily involved stakeholders. Following the discussions with a blind colleague, which gave us first ideas about the problems and needs of the visually impaired, and after we wrote as a result the first version of the scenario, we further sought the dialogue with a larger group of blind and visually impaired of different ages and education degrees. By that, we could gain a broader and at the same time deeper understanding of the problems and needs of this group, and were able to prioritize their needs as well. In the subsequent interdisciplinary dialogue, we discussed, among others, the economic cost and the acceptance constraints of the ideas from the first scenario. By combining the most urgent needs of the blind with the

---

[1] See [30], pages 93–107.

existing economic constraints, we could thus develop a prototype for the most realistic and useful part of the scenario. According to the scenario-driven prototyping method, we presented and discussed the prototype thus developed with both blind and visually impaired persons as well as with researchers from different areas. From the analysis of scenario and prototype we gained an abundant collection of both technological and societal insights.

The remainder of this chapter presents the scenario and the prototype of the Chatty Environment, as well as the technological and societal insights gained from the project, as follows: section 7.1 introduces the everyday problems of the blind and visually impaired in greater detail. Section 7.2 summarizes the Ladenburg "independent living" scenario, together with the results of the first round of discussions with blind and visually impaired. The prototype resulting from these analyzes is presented in section 7.3. Section 7.4 presents the technological and societal insights gained from the subsequent feedback given by stakeholders after having played with the prototype, as well as the multidisciplinary analysis of both scenario and prototype. Section 7.5 puts our work into the context of other ubiquitous computing assistive technology projects.

7.1 Motivation: Everyday Problems of the Blind and Visually Impaired

The blind and visually impaired encounter many difficulties in their everyday lives. Some of these problems, like their difficulty to find the way through the city, an airport terminal, or a large university building, seem to be easily understandable by sighted people. This is also the range of problems which can be partially alleviated with today's means. Guidance dogs are trained to know the paths typically taken by the blind – from home to work, to the local supermarket, or to the train station, for example. In many cities, pedestrian traffic lights emit a specific sound to signal the blind when green. Public places such as sidewalks, railway stations, or airports, are equipped with elevated guidance lines on the floor, which can be sensed by blind persons with the cane and used to maintain direction. As soon as the blind leaves his local environment, however, all these aids existing nowadays cease to be useful. In an unknown city the guidance dog cannot help with orientation. The guidance stripes along the sidewalk can only help the blind to remain on the sidewalk, not to find his destination. Likewise, the guidance lines inside the station of a foreign city cannot tell him at which platform his train leaves or where that platform is. Nor will they be able to guide him to the needed ticket vending machine, the restrooms, or to tell him whether the escalators run away from him, as desired, or come towards him. All these last examples point towards what numerous blind persons told us in interviews: As soon as they leave their domestic environment, they are reliant on external help. While they can (and do) lead a relatively independent life in their known surroundings, they cannot do so outside this geographic confinement.

Other problems are even less obvious, but affect them in leading an independent life as well. Sighted researchers are likely to learn of these problems only when speaking to the blind. Take for example the shopping in the local supermarket. Thousands of items, feeling all the same, spread over dozens of shelves, all the same shape. A blind person shopping on her own has thus great trouble finding the needed items. Since all packed food feels similar, the blind cannot use the tactile sense which helps her so well in other instances. Without external help, the blind person will only go to the known local supermarket and only buy a few items in learned locations. Furthermore, such everyday actions that we take for granted, like checking the price, the ingredients, or the best-before date of a product, are unthinkable for the blind. Another problem most sighted people are unaware of is that the visually impaired will often not be able to catch a bus because of its brief stop at the station, which is too short to allow her to find the door and the button to be pushed for opening it. Here again, blind people have to rely on external help.

Figure 7.1: Several parallel means of visually announcing a departing train at Zurich's main station. No means other than visual are provided.

Such difficulties seem to have a common cause – the lack of information about their surroundings that the blind experience. It is not only the shape of the surroundings, but also that much of the information we perceive is written. To emphasize this point, Fig. 7.1 presents several means by which a departing train is visually announced at Zurich's main station, while no other means than visual are provided. Why is this so? The cause probably lies in the nature of how humans use their senses to perceive the world – most people, when asked, will identify sight as the most important sense. This subjective impression is supported by anatomical facts. The eyes are a highly evolved mechanism, with about 150 million receptors. The human ear, for example, has only approximately 3000 receptors. More importantly, the brain region processing the visual input is with about 10 billion neurons more than five times larger than the brain regions handling any other sensorial input. Sight being the most evolved of human senses, the modern world is tailored to meet this fact, which worsens the problem for visually impaired. When constructing busses with buttons for opening the doors, it is likely that nobody thought of blind people and the trouble they will have finding those buttons.

7.2 The Ladenburg "Independent Living" Scenario

Bearing the problems encountered by the blind and visually impaired in mind, we have thus proposed in the "Independent Living" scenario the paradigm of a *Chatty Environment.* The main feature of the Chatty Environment is to enhance the visual information existing around us by other means of information that can be experienced by the visually impaired, such as spoken information. While moving through the Chatty Environment, this spoken information is continuously presented to the visually impaired user. Thus, he finds out how his surroundings are shaped and which entities exist around him, e.g., where the incoming bus goes to and where its nearest door is located, which package of food he is holding in his hand in the supermarket, or where the next fast-food-restaurant is located. The visually impaired is also able to get more in-depth information on selected parts of the environment and may even perform actions on some of these entities.

To realize this, the scenario envisions the usage of ubiquitous computing technologies to enhance the environment's real-world objects with a virtual component, which holds information about the corresponding object. To this respect, numerous real objects would possess an electronic beacon, which would create a *virtual aura* around the object, as

Fig. 7.2 exemplifies for the main hall of a railway station. When the user moves into the aura of such an enhanced real-world entity (or when the entity moves towards the person, as in the case of a bus), a mobile device carried by the user, which can communicate with the beacons, informs her about the object's existence and offers her a standardized interface for interacting with it.

Figure 7.2: The virtual auras of several objects of interest in the main hall of a train station.

7.2.1 World Disclosure and Navigation

The Chatty Environment provides two main functionalities to the blind or visually impaired user: *world disclosure* and *navigation*. With the first feature, the user's perception of the surroundings is enhanced by announcing him which entities he is passing by. Having the environmental entities in the immediate neighborhood being announced to him, the user acquires an extension of his own world perception. With the second feature, the Chatty Environment can also navigate the user to different points-of-interest.

Our original scenario presents the combined advantages delivered by these two functionalities to a blind researcher from Hamburg, who comes to visit a professor at the ETH Zurich. In today's scenario, the researcher encounters such a large amount of difficulties that his visit would be almost impossible without a sighted companion: Since all airports and train stations (but for the Hamburg station) are unknown to him, he will not only have great trouble finding the ticket vending machines, check-in counters, restrooms, money exchange boot, or the fast-food corner, but even the gate to board the flight or the right track for the needed local trains. Due to the earlier presented problems of a blind when shopping in an unknown environment, the scenario's protagonist can also not bring back home to his girlfriend the Swiss chocolate brand that she so much likes.

The future scenario shows to which extent these problems could be alleviated through the above-mentioned two main features of the Chatty Environment, world disclosure and navigation. Upon arriving at the unfamiliar Hamburg airport, his "smart cane" guides the blind to the air company's check-in counter, being aware of which company he flies with after having communicated with the electronic ticket. If there is only a short time left until boarding, the next default action would be to guide him directly to the gate. If there is enough time, the system computes and presents to the user a list of all possible pasttimes, such as airline lounge, restaurants, or shopping opportunities. There is no artificial-intelligence-like functionality in this feature (or, as Lueg [106] puts it, no approximation of the user's wishes), since the system does not make any choices; from the distances and time information it merely computes an unsorted list of possibilities. Since, however, there are hundreds or thousands of possible destinations inside an airport, the scenario postulates following rule: While a potentially unlimited number of objects might be electronically enhanced with a virtual aura, some targets inside a complex building such as an airport terminal or a railway station have a distinctive attribute. These entities, which are the same that are announced on the physical signboards for sighted travelers, will be the only ones taken into consideration when presenting the user the choice of pasttimes. With this feature, we have tried to map into the Chatty Environment the way sighted people are being presented with unknown environments.[2] If he wants to be guided to a destination not contained in this predefined set, the user will have – as any sighted user, too – to choose the information booth, let the system guide him there, and receive there new directions. More than a sighted traveler, though, he also has the option of entering the wished

---

[2] As will be presented shortly, there is one more concept by which we have tried to map into the system the way sighted people perceive the world: the extent of the virtual aura.

destination into his mobile device communicating with the Chatty Environment; if the destination has been electronically enhanced, he will be guided directly there.

On the other hand, the blind user can also discover features of the environment (e.g., the airport terminal) while passing by, in precisely the same manner as the sighted traveler does. This feature is the very consequence of the world disclosing property of the Chatty Environment. When passing by the restrooms, for example, these are announced to the user as all the other electronically enhanced objects, too. By explicitly allowing to adjust the radius of an object's virtual aura, we have provided a second tool to map into the Chatty Environment the way that sighted people discover unknown environments. Large objects such as the railway station itself should – according to our paradigm – emit a stronger signal, so as to be perceived by the blind user from a larger distance than smaller or less important objects.

After successfully having boarded the plane, the researcher lands at Zurich airport and takes the train to Zurich. The Chatty Environment not only guides him to the right track, but also offers the means for him to undertake actions in the real world, e.g., to remotely open the train's door without having to search for the button that has usually to be pressed for this action. In Zurich, he still has time to shop for his girlfriend's favorite Swiss chocolate in the station's supermarket. Once in the supermarket, he enters the desired brand into his mobile device, which communicates this to the supermarket indoor localization and navigation system. The supermarket guides him to the shelf with chocolates and offers the information he requests about the products' contents and expiration dates. While walking towards the chocolate shelf, however, he walks by a board announcing a special offer for fondue sets. Since the board is electronically enhanced, he is being announced the available special offer and ends up buying a fondue set as well. On time and with his shopping done, he then walks to the ETH for his meeting.

7.2.2 Analysis of the Scenario and Priorities of the Visually Impaired

The scenario uses a combination of world disclosure and navigation to provide the blind with the possibility of independent traveling. While there are several instances within the scenario where the world disclosure feature is used, such as the special offers in the supermarket or the hint towards the restrooms at the airport, the main focus of the "independent living" scenario lies on the navigational part. This priority came during the dialogues within the Ladenburg collegium, and represents the opinion of the participating scientists that navigation should obviously be the more important need of the blind in order to be able to lead a more independent life. When we presented the scenario to the blind themselves, however, they drew a somewhat different picture from what we had expected, as will be subsequently presented.

In order to deepen our understanding of the problems and wishes of the blind and visually impaired, we have contacted the Swiss Association of the Blind,[3] which has pointed us towards members who would potentially be interested in providing us insights into their needs and test our prototype when ready. Our questions for them revolved around following main topics:

- General questions about the Chatty Environment: How would they feel about objects that "talk" to them? Which is the more important feature of the two, environment disclosure or navigation? From the technological possibilities we present, what would be valuable for them, and what not?
- Navigation details: How should a navigational aid for the blind and visually impaired best describe a path through an unknown environment?
- Usability: What sorts of output are preferred and should be thus used by the mobile device? What ergonomic issues have to be considered? How could the new means replace or how could they meaningfully complement existing aids such as the blind cane or the dog?

We were able to conduct interviews with nine blind and visually impaired, five women and four men. The questioned people were at the time of our interviews between 30 and 81 years old, with a medium age of 54. They live in different regions of Switzerland and their educational level varies from high-school without degree to university degree. Two of our interviewees (aged 79 and 76, respectively) are a married couple. Table 7.1 presents this data in an overview.

---

[3] "Schweizerischer Blinden- und Sehbehindertenverband," see www.sbv-fsa.ch/.

Person	Gender	Age	Education	Residence
Person A	male	57	high school degree	Zürich
Person B	male	33	university	La Plaine
Person C	female	30	computer school	Corcelles
Person D	female	34	high school	Neftenbach
Person E	female	81	high school	Winterthur
Person F	male	79	university	Wiesendangen
Person G	female	76	teacher school	Wiesendangen
Person H	female	49	domestic school	Winterthur
Person I	male	55	university	Solothurn

Table 7.1: Gender, age, education, and place of residence of the interviewed blind and visually impaired.

The impairments of the interviewed range from total blindness to 40% of sight. Most completely blind are blind since their birth. For the others, the time since the impairment started varies between 1.5 years and more than 20 years. Some of them are professionally active, some work with computers. All completely blind were in a specialized school for blind and all except one person use the regular blind cane or its shorter version. Table 7.2 presents this information in detail.

Person	Blindness	Since	School	Walking Aid
Person A	blind	birth	yes	cane
Person B	blind	birth	yes	cane
Person C	blind	birth	yes	cane
Person D	97%	5 y	yes	cane, dog
Person E	86%	10 y	no	short cane
Person F	60%	21 y	no	none
Person G	85%	1.5 y	no	short cane
Person H	~100%	13 y	no	cane
Person I	blind	20 y	yes	cane

Table 7.2: Blindness degree, period of impairment, profession, special school attended, and used walking aids of the interviewed blind and visually impaired.

The interviews were conducted in two steps: All interviewed first answered a questionnaire comprising 20 questions, ranging from general information about their age, profession or impairment grade to precise questions about use of handheld devices, preferred input and output methods, and particular requirements for object descriptions. The above-mentioned three main topics were either represented through several questions or by one question needing a long and detailed answer. To be able to understand how blind and visually impaired would like to be navigated through unknown environments, for example, one of the questions required each interviewee to describe a known path in a train station, as he or she would describe it for another visually impaired person, who does not know the respective station. The interviews were further based on the "open-end" principle, each participant being able to add any information or suggestion considered to be relevant. The interviews were typically between one and one and a half hours long. The answers can be grouped into four categories:

Chatty Environment

The environment endlessly speaking to the user, telling her about the surroundings, may seem an annoying property to most sighted people. However, all interviewed said they would probably not be disturbed by objects that "speak" to them. They rather experience a substantial lack of information today and welcome any additional information, especially the completely blind users. They imagine that these objects would help them to find their way more easily, even without an explicit guidance aid. Six out of the nine participants, however, when explicitly asked, said that too much information could ultimately become annoying and make them feel uncomfortable. Thus, they would prefer to be able to adjust the number of objects that speak, i.e., the information density, depending on their actual needs.

As a surprise came for us the relative weight that users gave to the two foreseen aids of the system: environment disclosure and navigation. We had asked the question about their importance being almost certain that our impression that the navigational part is more important will be confirmed. We had foreseen navigational aids as being the more relevant; environment disclosure, while possessing a value of its own, was mainly depicted as an aid for the blind and visually impaired to know which entities exist in their surroundings, select them, and being then navigated towards them.

During the interviews, however, and especially during the second part consisting of a free discussion, the large part of the interviewed either explicitly named the world disclosure functionality as the more important, or were much more interested in discussing the properties of this part of the scenario. They asked numerous questions about the exact features that would technologically be feasible, and told us about many problems they encounter in their daily lives, and which could partly be solved by an environment disclosure property, even without a navigational part.

Many of the everyday problems of the blind, as described in the motivational section 7.1 of this chapter, have been revealed to us during this part of the interviews. We could learn about problems experienced by the blind while shopping or when trying to use the public transportation system, for example, which are all related to the lack of access to information. It is not only the perhaps more sophisticated information about the ingredients of a product or its expiration date which is missing, but even most basic information such as which product the blind person is holding in her hand in the supermarket or which bus is coming to a halt at the stop in front of her. All this information is unrelated to any navigational issues. It is rather about providing the blind or visually impaired user with contextual information about her environments, or, in other words, what we have called "world disclosure."

Related to the public transport discussions with the blind, we could also learn about the difficulty they experience in finding the buttons used to open the doors of busses or trains. When distillating the findings from the interviews, we realized thus that using the Chatty Environment infrastructure and the blind user's mobile device to enable him to not only gain information from the real world, but also use the back channel to electronically trigger actions in the real world, could be a valuable complement to the original ideas.

Path Description

Despite the preference for the environment disclosure part, we also received an important amount of feedback related to the way visually impaired would like to have a path described to them by a navigation system. Visually impaired describe their routes with much more details than sighted people. Completely blind tend to be even more detailed than people with partial sight. In particular, the blind often use objects such as walls, balustrades, or murals in order to define directions. Other important objects used in these descriptions are places with a special sound or smell, such as fountains or bakeries. Further important information can be obtained from the configuration of the floor such as the end of a pavement, or the change in its texture. The individual insights about the system's path calculation and presentation are presented below.

The calculated path should avoid elevators. All interviewed persons have problems with elevators as they are often not able to find the right button to press. Elevators should only be used if the system knows that they offer blind support, i.e., buttons with inscriptions in Braille or relief letters. The system should also avoid escalators and prefer stairs instead. While visually impaired do not have particular problems with escalators, they are typically unable to use them when escorted by guiding dogs. The coordination between dogs and their owners is difficult to achieve, since the dog typically walks in front and would drag its owner forward once it has reached the escalator. Likewise, ramps should be preferred to stairs. Ramps are the most natural way to go up or down, as they are regular, without any edges or interruptions.

The system should try to use a *rectangular* approach, i.e., the path should consist by a series of straight segments that always intersect under a 90-degrees angle. If this approach is not possible, the “clock navigation” is a good way of indicating angles. Blind can distinguish well between *two o’clock* and *ten o’clock*, while instructions such as “40 degrees to the left” or “35 degrees to the right” would cause confusions. The path should further run along walls rather than in the middle of open spaces. Visually impaired, and especially blind, can experience insecurity when in the middle of a large space such as the entrance hall of a large railway station. They are further more unlikely to use easily recognizable smells or sounds in the middle of such space than along its borders.

The path should be divided in series of segments not longer than 5 to 10 meters, even if these segments belong to the same straight line. Visually impaired would like to have regular feedback on their current position. Giving a new instruction every 10 meters assures them that they are still on the right path. The path should end at the door of a shop (or bus) and not just somewhere in front of it. One of the major problems for the blind is finding the door, once they are in front of their destination.

Usability

Only two interviewed persons own a Personal Digital Assistant (PDA) for the blind. Nevertheless, most of them had had such a device in their hands and have an idea about their approximate weight and measures.

Speech is the preferred output medium for almost all interviewed persons. Some of them could find utility in additional signalling techniques, such as vibration or non-speech audio signals (i.e., beep). All rejected the idea of force feedback on the blind cane, as the normal use of the cane would be altered too much through this technique. Speech output should be understood without major difficulties; the speed should not be too high, the spoken instructions have to be clear and well articulated.

All participants would like to carry the device in the pocket (or handbag for women) or hanged around the neck, in order to keep the hands free for other functions. One hand is almost continuously occupied with the cane or the dog's leash, the mobile device should not occupy the other hand as well. All interviewees rejected the idea that an electronic system could substitute today's means of orientation, especially the cane that enhances the tactile feedback, so important for the blind. Novel means can only be seen as an informational and perhaps navigational complement to today's proven means.

The device being stored away in a pocket, its acoustic output needs to be transmitted to the user via a headset-device. Nevertheless, important about speech output (and audio output in general) is the fact that the normal hearing of the user must not be altered by this system. Blind need stereometric hearing in order to determine for example the direction of moving obstacles. Therefore, any kind of headphones or earphones used has to comply with this requirement. This excludes stereo headphones; mono headphones are suitable only if they let environmental sounds get through.

Additional Issues

Some additional information could also be retrieved from the interviews. Security is an important issue. Almost all interviewed persons prefer a longer (and consequently less quick) but safe route to a shorter but more risky path. The system should further be able to determine when the user misses the path and autonomously notice him if this happens. It should stop him and indicate how to reach the destination from the current position.

The quantity of noise is also an important point: Instructions should be more detailed and more frequent in noisy environments, as visually impaired often feel anxious when passing by regions with too much noise. These regions should be avoided when possible.

7.3 The Prototype: Chatty Environment and World Explorer

After evaluating the interviews with blind and visually impaired persons that have been presented in the last section, we started to build the Chatty Environment's first prototype. Due to the then-surprising new insights into the priorities of the blind, as well as the fact that several other projects were already exploring navigational aids for the blind (as section 7.5 will show), while almost none was investigating environment disclosure properties, we decided to focus our efforts on this latter functionality.

Due to its main focus, the resulting prototype has been named "World Explorer." The current section is dedicated to present the several development stages of the prototype, which we will address alternatively as "World Explorer" or "Chatty Environment," depending on whether the discussion focuses more on the mobile device itself, or the whole infrastructure, respectively.

7.3.1 First Chatty Environment Prototype – Components

We implemented the environment disclosure property of the Chatty Environment by using pervasive computing technologies to enhance the environment's real-world objects with a virtual component that holds information about the corresponding object. For doing so, real-world objects were tagged with an electronic device that either directly holds the information, or points towards other data sources containing this information. Since the tag continuously sends this information out to the world, similarly to a lighthouse on the shore, in the original works [33, 135] we have thus often referred to the tags as "beacons." In this section, the two terms are used interchangeably.

The beacons create the virtual aura around the object they tag, as described above. When the user moves into the aura of such an enhanced real-world entity (or when the entity moves towards the person, as in the case of a bus), the device carried by the user – the World Explorer – informs her about the object's existence and offers her a standardized interface for interacting with it. This first prototype has been extensively described in the diploma thesis of Felix Röthenbacher [135][4] and highlighted in [33]. We summarize below the most important features.

---

[4] The majority of figures used in this section also originates in Felix' work.

Tagged Entities Spread in the Environment

All objects of the Chatty Environment are marked with small electronic radio devices. These beacons generate the objects' auras, thus making them detectable, as stated in section 7.2. For the first prototype, we have used the so-called *Berkeley motes* as beacon devices that signal to the world the presence of the object they tag.

The motes are among the first and still best-known purpose-built nodes for sensor networks. They have originally been developed during the late 1990s at the University of California, Berkeley, at the universities' "sensor and actuator center."[5] Their development is now commercially being pursued within the spin-off company "Crossbow."[6] The first generation of motes, called "WeC" [73], appeared in 1999. For our prototype, we used the "Mica2" and "Mica2Dot" motes [126] (depicted in Fig.7.3), launched in 2002. As with all Berkeley motes, Mica2 and Mica2Dot also run the TinyOS operating system,[7] jointly designed with the motes and tailored to meet the needs of sensor networks applications.[8] Applications for the Berkeley motes are being written using the "nesC" (network embedded systems C) [59] programming language, a dialect of C with support for the structuring concepts and execution model of TinyOS.

For the first version of our prototype, the motes seemed the most appropriate platform to be used as beacon devices for several reasons: Firstly, being conceived as nodes for sensor networks, they can easily be extended to contain sensors. The ability to sense parts of their environment, and thus provide the visually impaired user not only with the static data stored in the information file, but also with dynamic, up-to-date information, seemed to be a valuable feature for the system, although we had no immediate usage for it. For the future, however, we envisioned usage scenarios like a bus which continuously changes its position. If the current position should be transmitted along with the other data about the bus (or other information, derived from the position, such as the next stop), the possibility to attach a GPS sensor to the tagging device is valuable. Secondly, as opposed to less powerful devices such as RFID tags, for example, the motes have a decent

---

[5] See www-bsac.eecs.berkeley.edu.

[6] See www.xbow.com.

[7] See www.tinyos.net.

[8] TinyOS has later been extended to work with other sensor networks platforms as well. One example are ETH Zurich's BTnodes, which, alternatively to a dedicated operating system can also run TinyOS. See also www.btnode.ethz.ch/Documentation/TinyOSOnBTnodes.

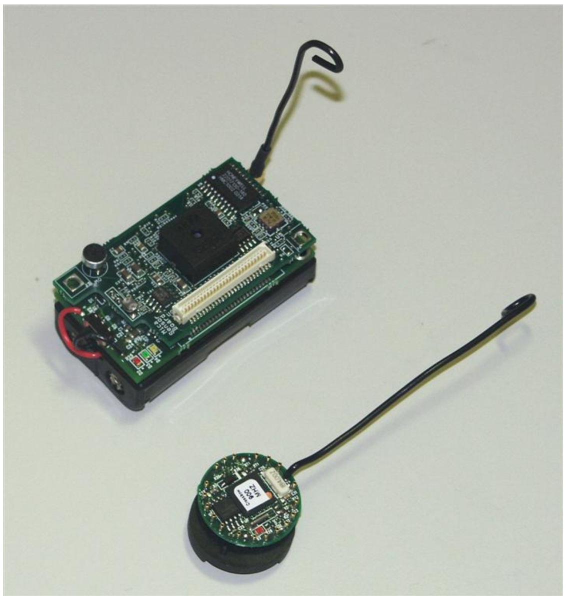

Figure 7.3: The Mica2 (above) and Mica2Dot (below) sensor nodes used as tagging devices for the first prototype of the Chatty Environment.

amount of storage space, enough to store the required information file about a real-world object. By being also equipped with an autonomous power source (two AA batteries), they have a maximum transmission range in the order of 50-100 meters, which is roughly the amount we needed for our scenario. A noticeably smaller range would render much of the scenario not practicable. Thirdly, and for us the decisive argument, has been the fact that the motes offer the advantage of adjustable emitting power. Their transmission range, and thus the virtual aura's radius, can be fine-granularly tuned from a small range of about 5 centimeters to the maximum range of about 100 meters. This property of adjustable range enables us to create large auras around large or important objects, and smaller auras around less important or small objects, recreating the way sighted people would gather information about those objects.

One of the weaknesses of the motes, on the other hand, is the proprietary radio protocol they use for wireless communication. As opposed to other sensor networks nodes, such as ETH's BTnode,[9] which uses the industry standard Bluetooth radio communication, motes cannot easily communicate with other devices. To circumvent this problem, we had to connect a mote to the serial port of the user's mobile device, as a communication interface to the other motes spread in the environment.

---

[9] See www.btnode.ethz.ch.

Mainly because of this last drawback, we have used a modular, easily extendable design, to allow for further tagging devices to be added to the system in the future. Using other beacons is a relatively easy task, as section 7.3.2 will show.

World Explorer Device

The first World Explorer device (see Fig. 7.4), carried by the user, is both beacon reader and mobile computing platform. It constantly sends identification requests. If a beacon receives this request, it sends a reply with a unique identification. Upon detection of a beacon, the system creates a beacon object and sets it to an active state. Beacons have to store a unique identification and a human-understandable textual description.

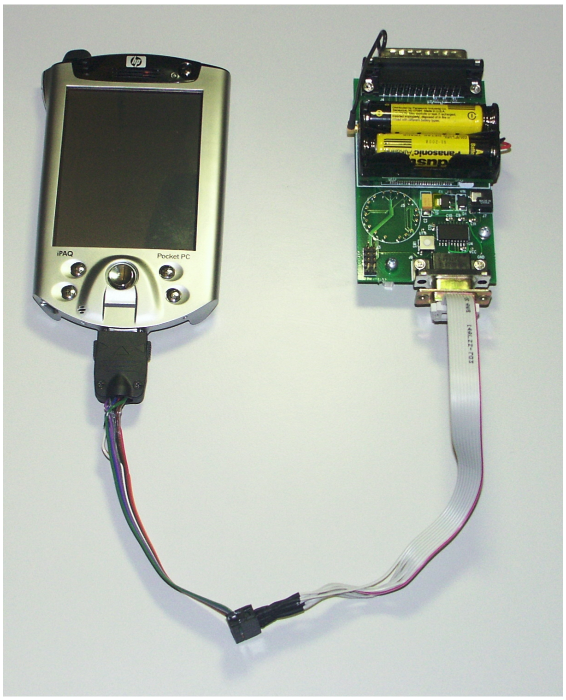

Figure 7.4: View of the first prototypical World Explorer device.

As World Explorer device, we used an HP iPaq 5450, with integrated WiFi 802.11 and Bluetooth connectivity. At its serial interface, a mote is connected, which sends out the inquiries for motes in the environment and transmits the answers to the software running on the mobile platform. Aside from ID and textual description, further information may be stored on the beacon, but also on its virtual counterpart, as will be discussed below.

Virtual Counterparts

While the motes used for the first version of the prototype could store in their 4KB of RAM memory enough textual information about the object they tag for all practical purposes, we envisioned from the project's outset to extend the World Explorer to communicate with other tagging devices as well. Not only would this be needed for building a flexible and general-purpose system, which is an important feature for a system designed for the relatively small group of visually impaired.[10] This feature was also needed to allow the prototype to cover situations characterized by a large number of small items, such as supermarket products. As section 7.3.3 will show, to include this part of the scenario into the prototype, a different interaction paradigm with tagged items in the environment was required, a paradigm that could most easily be offered through RFID tags. Furthermore, passive RFID tags are about two orders of magnitude cheaper and also much smaller than sensor nodes, making it thus likely that in a near future numerous objects (especially products in supermarkets) will be tagged by them, while sensor nodes will most likely remain sparsely distributed.

Passive RFID tags, however, usually have no RAM memory to store an information file about the object they tag. They only store a unique ID in a read-only memory. The information about the object has thus to be retrieved from another location, using the ID communicated by the tag as access key. The concept of digital representations of real-world objects, representations that store information about these objects and are used by applications to interact with, has been introduced in [94] and then in [132] under the name of "virtual counterparts." [132] states that: "Due to resource limitations, neither the physical object nor the tag is able to implement all of the above concepts. Therefore, a digital representation is needed – the virtual counterpart of a tagged object – that can adopt this role." A first thorough analysis and implementation of virtual counterparts is presented in [45].

We have modified the concept of a virtual counterpart to allow for several data sources to contain the information about a real-world entity. In our definition, one of these data sources might be located on the tag itself (together with the ID), others at remote locations. The only condition is that the links to these remote locations are stored

---

[10] Since this group is too small to attract large investments, a system designed for it should be able to make use of technology that has already been deployed into products for other reasons, as our multidisciplinary dialogue has shown [30].

on the tag itself. This also means that the virtual counterpart does not exist as an entity in any of the remote data sources. Parts of it are rather scattered across several locations, and the virtual counterpart can only be composed in the World Explorer mobile device itself. Upon discovering a new object, the World Explorer gathers the information transmitted by the tag itself (if any), joins it together with all information from remote sources, and only then presents the object to the user. For RFID tags, however, there will typically be no data source on the object itself, and only one remote data source.

A further related concept used in the World Explorer system is that of a "virtual beacon," which only partly overlaps with the "virtual location" from [45]. A virtual beacon does not correspond to a physical object, but only exists inside the virtual world of the Chatty Environment. Obviously, the tag announcing the virtual beacon has to be physically placed somewhere in the real world. The real-world object it is placed on, however, is not the one it logically corresponds to. Such virtual beacons will typically be used to announce information about an object that will soon come close or an event that will soon occur, and will be relevant for the users in a given context, but which is still either too far or is a purely abstract concept. Röthenbacher [135] underlines the concept by the following example: "A virtual beacon can be used in a train by the engine driver to announce the next stop. A few minutes before arriving there, a virtual beacon with information about the stop is activated. The user is informed in the usual way that a beacon has been detected, although the beacon exists only in the World Explorer system. He has the possibility to browse the beacon to get information about the name, connecting trains, etc" [135].

7.3.2 World Explorer Software Design

From the concepts described above, and as shown in Fig. 7.5, the overall system consists of the blind user's mobile device running the World Explorer software, the beacons tagging objects in the environment, and, possibly, external data sources, which – as described in the last section – may also contribute with data for the generation of virtual counterparts. Fig. 7.5 further shows the communication protocols between these individual components.

The World Explorer software has been constructed from the outset as to be as flexible and easily extensible as possible. To this respect,

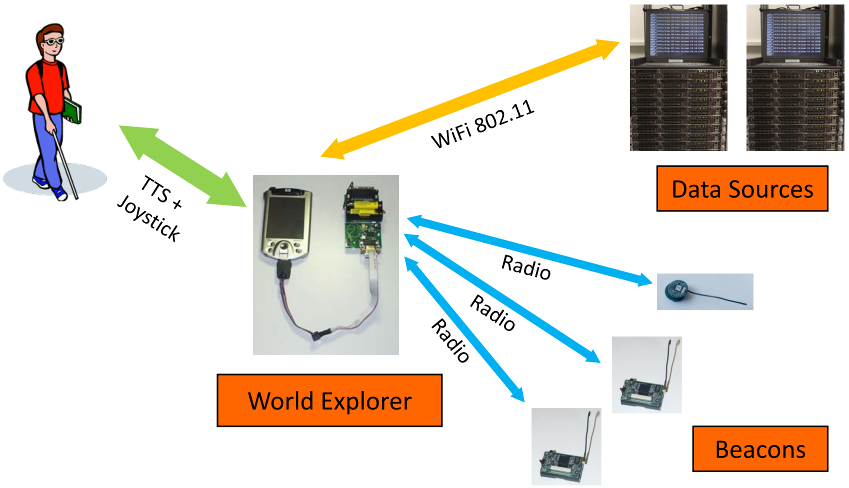

Figure 7.5: Top-level view of the Chatty Environment system including the World Explorer software on the user's mobile device, the information beacons tagging real-world objects, and the external data sources.

the component-oriented software construction paradigm [154] has been used, which allows the building of applications from independently developed binary components, without the need of recompiling already existing code. Due to the features offered by the mobile device, we decided to write the World Explorer software in C++ using Microsoft's eMbedded Visual C++. As component model, it has thus been a natural choice to opt for Microsoft's Component Object Model (COM) [169] as well. Before describing the component system, we will first describe the communication protocols and the data model used for the prototype.

Communication Protocols

As shown in Fig. 7.5, the first version of the World Explorer uses two communication protocols:[11] the proprietary motes radio communication for receiving data from the beacons, and http requests sent over WiFi to the external data sources. First, the mote "base station" (i.e., the mote connected to the serial port of the mobile device) is used to detect beacons in the environment. To accomplish this, the base station is continuously polling the environment. If a beacon is found that was unknown so far, a beacon object is created. The object is initial-

---

[11] The interaction with the user consists, as the picture depicts, of spoken information about the environment (generated by a text-to-speech engine), and of user choices through the device's joystick. These features will be introduced shortly.

ized with the data that originated from the beacon. Then, all available information sources (linked by the beacon) are enquired through WiFi and the data is added to the object's state. Only after joining all this data, the newly found object is presented to the user.

As will be shown shortly, the proprietary radio communication is encapsulated inside a "beacon system" component. Adding new sorts of tagging devices (such as RFID tags) to the system will thus only consist in writing a new component that encapsulates the communication with the new tags, and adding this new component to the system. This extension will not influence already existing communication components in any way.

Data Model

The standard data format used throughout the World Explorer system is the Extensible Markup Language (XML).[12] The external data sources provide XML information files. The information files stored on the motes are in the same format. The communication between the system's components consists of XML files. The final information presented to the user is extracted from an aggregated XML file, containing all available digital information about a real-world object. The only exception to this rule are proprietary data formats, such as the IDs of RFID tags. Such exceptions, however, have to be handled inside a communication component, which then presents to the rest of the system information in the XML format as well.

Aside from the thus-achieved standardization, since text files are much smaller than audio-files containing the same information two more advantages are achieved by this approach: First, more information can be stored on the beacon itself. With devices like the Berkeley motes, having a reasonably large memory, it becomes thus possible to store all the needed information on the beacon itself. Saving the effort to access remote data sources improves the flexibility and the robustness of the system. Second, the data exchange is done more rapidly, thus giving the user an up-to-date view of the environment.

Last but not least, the definition of a XML schema[13] bypasses compatibility issues between different platforms and opens up the system for third-party information providers. For the Chatty Environment, we thus defined the XML schema shown in Fig. 7.6. The most important

---

[12] See www.w3.org/XML/.

[13] See www.w3.org/TR/xmlschema-0/.

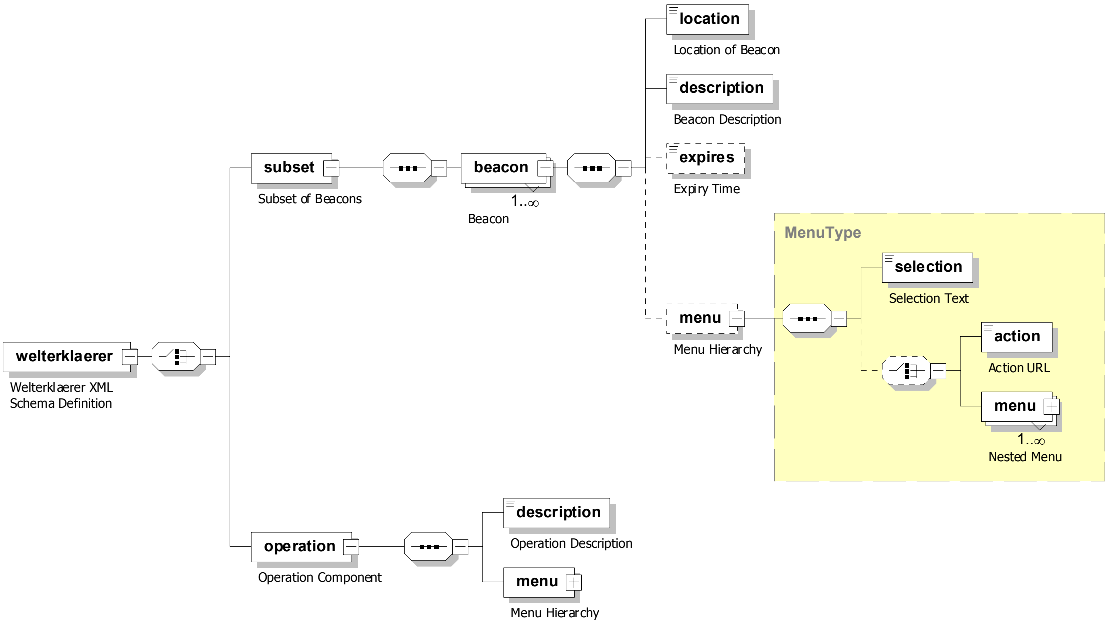

Figure 7.6: Chatty Environment XML schema.

part is the highlighted nested hierarchy defined by the `menu` tag. Each menu entry has a selection text which is presented to the user. There are three types of menus: information menu, selection menu, and action menu. If a selection or an action menu is selected, either the submenu is presented to the user, or an action is executed, respectively.

System Components

Any of the system's software components has to implement one of the five types of interfaces presented in Fig. 7.7 (`ComponentSystem` being the root interface):

- `InputSystem` for components that allow the user to provide input (for example, by selecting one of the presented objects to learn more information about).
- `IOSystem` for components that provide information and real-world interaction possibilities to the user, such as triggering an action.
- `BeaconSystem` for components that can discover a specific sort of tags and encapsulate all communication with these. The communication protocol might be proprietary.
- `InformationSystem` for components that gather information from external data sources (i.e., information not stored on the tagging

device itself).

- `ActionSystem` for components that will implement the execution of real-world actions, such as opening a train's door. For this example, the (proprietary) communication with the electronic device on the waggon's door will be encapsulated in such a component.

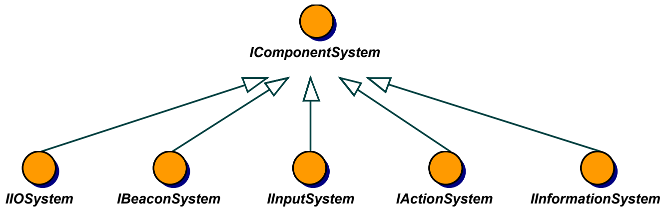

Figure 7.7: The five types of software components of the World Explorer: IOSystem, BeaconSystem, InputSystem, ActionSystem, and InformationSystem.

For the first version of the World Explorer, the components depicted in Fig. 7.8 have been implemented. They all belong to one of the five types described above, except for the component manager, which manages the whole application.

The component manager handles the different system components, which belong to one of the five different component categories. For every component category, it maintains a list with all momentarily available components. If a component is needed, the component manager is asked for an enumeration of the available components. The component manager further assures that every component is initialized before use.

For this prototype, a single `InputSystem` component exists. It handles the user's input from the five-way joystick of the mobile device, as depicted in Fig. 7.9. Two `IOSystem` components for the interaction with the user exist: a text-based, and an audio-based. The text output will obviously be of no use for the blind user. It has been implemented for testing purposes by sighted researchers and also as a proof of concept for further applications which may deploy the World Explorer, such as a tourist guide for sighted tourists. The audio output system will be presented shortly, together with the joystick input system.

Two `BeaconSystems` have been implemented. The `motes` component retrieves the information stored on the motes themselves and handles the proprietary communication with them. The `beaconclient` component retrieves data from a software simulation of beacons, running

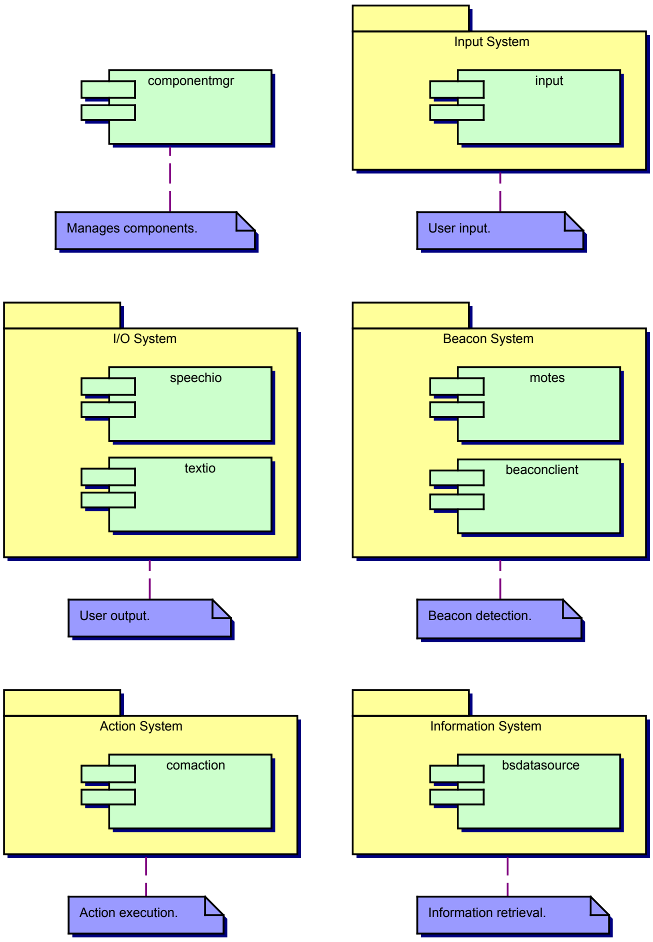

Figure 7.8: Software components of the first World Explorer prototype.

in our lab for testing purposes as well, the `bsdatasource` component. This server is a small network application which allows through the simple interface presented in Fig. 7.10 the simulation of arriving and disappearing beacons. A simple user interface is provided for loading, activating and deactivating specific beacons. When the beacon client is active, it continuously polls the server for state changes and notifies registered listeners about detected and lost beacons, while the server is listening on port 8092 for such requests. Since for this first prototype, there are no further tagging systems than the motes, and the motes store all information about the object they tag themselves, there is no further `InformationSystem` beyond this simulation of beacons.

Finally, the system prototypically implements a simple action system, as a proof for the concept of triggering actions in the real world. We have therefore attached a sensor board to the motes that, among other sensors, possesses a sound device that produces a loud “beep” when triggered. Since for some scenarios it would be useful for the

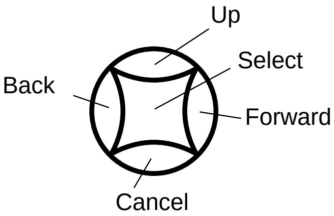

Figure 7.9: Input system for the first prototype, interpreting the input from the five-way joystick of the mobile device.

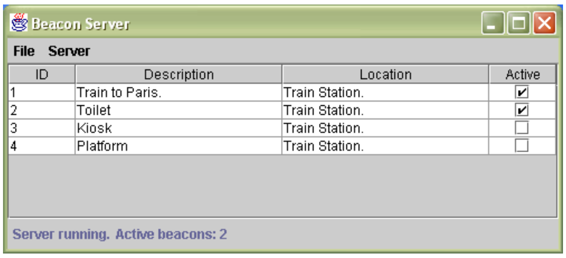

ID	Description	Location	Active
1	Train to Paris.	Train Station.	☑
2	Toilet	Train Station.	☑
3	Kiosk	Train Station.	☐
4	Platform	Train Station.	☐

Server running. Active beacons: 2

Figure 7.10: Interface to the beacon server, which allows an interactive simulation of appearing and disappearing beacons.

user to be able to precisely locate a close object through acoustical means, the Chatty Environment's objects augmented with motes offer the possibility of triggering this beeping sound. Obviously though, the true value of the `ActionSystem` concept lies in allowing more sophisticated actions, such as the communication with trains' doors and the triggering of actions such as opening them. For practical reasons we could not implement such functionality; it would, however, be an easy task with the provided framework of action components.

User Interface

Inside the `speechio` component, a lightweight third-party text-to-speech-engine,[14] supporting nine different languages, generates the spoken output needed by the blind or visually impaired user. The user interface is designed to allow the user to operate the device by responding to speech output. A five-way cursor serves as input device. The keys are conceived to resemble the possible interactions with a web browser, namely, *select, back, forward, up,* and *cancel*, as well as a dedicated *home* key (not depicted in Fig 7.9).

---

[14] Formerly, Elan Tempo PocketSpeech engine, a multilingual text-to-speech engine for PDAs. Now sold under the name "Acapela Mobility," see www.acapela-group.com/acapela-mobility-2-speech-solutions.html.

The information is organized as a logical tree that can be navigated by the user. The topmost level of the tree is represented by the objects active in the system at any given moment. The information for every object is then itself organized as a subtree. The subtree's root is the designation of the object. It can be freely chosen, but will usually be the object's textual description in a meaningful way for the user. The nodes represent choices, the leafs either contain a specific information, or the triggering of an action. The system starts at the topmost level, announcing the active objects (i.e., the ones in the neighborhood, whose signal can thus be received). They are sequentially presented to the user, with a pause of two seconds between them. To choose an object, the user has to press the *select* button while the object's name is read, or in the two second interval afterwards. The logical navigation within an object is then realized in the same manner. Thus, to receive for example the information about the time a train arrives at the wished destination, the user has to follow following path: "Train to Paris," *select*, "Departing time," "next stop," "all stops," *select*, "Berna," "Lausanne," "Geneva," *select*, "arrival time," *select*.

The usage of the other keys is straightforward. Especially relevant is the *back* key, typically used by the user if he did not press on time the *select* button. The *home* button brings the user back to the topmost tree level, where all existing objects are presented. The communication with the user is based on natural language. Each menu item consists either of an introductory text (e.g., "Please choose:"), followed by an enumeration of the subtopics, for a node, or of an information text (e.g., "The traffic light is red."), for a leaf. If the menu item is associated with an action, the action is executed when the user selects this item.

The World Explorer further supports two message queues with different priority levels to inform the user. The high-priority queue is used to announce warnings about obstacles, red traffic lights, etc. The normal priority queue is used to present ordinary information to the user. When a high-priority message arrives, the currently processed message is interrupted and the high-priority message is read at once.

7.3.3 New Interaction Paradigm through RFID Tags

While real-world objects that are tagged with motes present numerous advantages for several of the usage scenarios, they also have a set of drawbacks: the motes (or any other sensor nodes) are too large or too expensive to tag small or cheap objects, they rely on the energy supply from their batteries, and there is still no sensor node industry standard in sight. More importantly, though, for some of the usage scenarios, their large range, otherwise an advantage, can also represent a disadvantage.

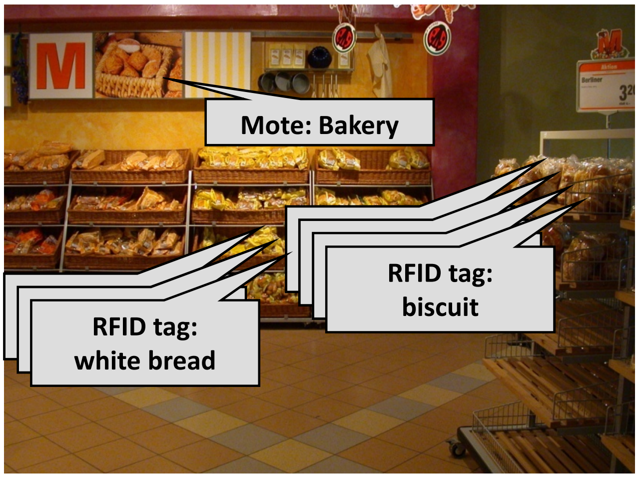

Figure 7.11: A supermarket bakery using a mote for macroscopic orientation, and RFID tags for the microscopic orientation.

Such a scenario is pictured in Fig. 7.11. It is certainly a good idea to have the individual sections inside a supermarket, such as the bakery in the example, tagged by motes so that they can announce themselves to the user. Aside of the financial impracticability, however, if every individual product was tagged by a mote as well, the user would be received by a flood of information upon entering the supermarket or one of its sections. Hundreds, maybe thousands, of partly identical products would announce their presence.

What the blind user rather expects in such an environment is to know where each sort of product is located, and, when there, to be able to explicitly select one or the other product, in order to compare ingredients or expiration dates, for example. For such a scenario, the very small communication range of passive RFID tags (for some systems as low as a few centimeters) becomes an advantage. The user has to explicitly select a product, and will receive information only about that precise object. For such small "neighborhoods," the list of active objects (i.e., "objects in the neighborhood") will comprise only one object.

As Fig. 7.11 shows, the two components, motes and RFIDs, thus complement each other in an ideal way for such a scenario. Additionally to the *macroscopic* view of the world provided by the sensor nodes used to tag large or important objects of the Chatty Environment, RFID tags provide a *microscopic* view that highlights specific details. Since the advantages of this complement seem evident, the first extension of our system has thus been the inclusion of RFID tags as "beacon" (microscopic beacons, that is) devices.

After comparing the features of several systems, we decided to use one of the simplest systems available, the μ-chip system from Hitachi.[15] The μ-chip, at that time the smallest commercially available RFID chip, possesses several features that made it well-suited for the Chatty Environment, such as the small size of reader and the antenna (as depicted in Fig. 7.12), or the small reading range (even compared to other RFID systems) of about 5 centimeters. The former attribute is important since both reader and antenna have to be carried by the user; the latter, since it matches our wish for a "microscopic" approach.

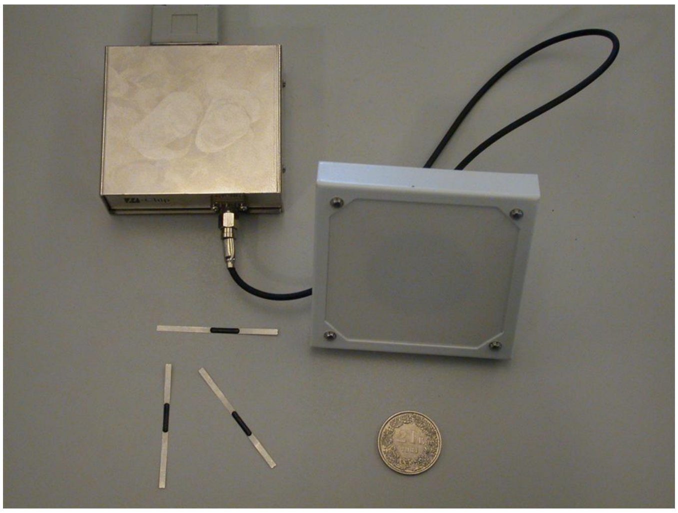

Figure 7.12: Hitachi μ-chip reader (on top), antenna (on the right), and tags (lower left corner), as compared to a two Swiss Francs coin.

The μ-chips possess a 128 bit non-exchangeable ID. They transmit on the frequency of 2.45 GHz and are thus inside the ISM band, which is freely usable in numerous countries. They do not have a collision avoidance protocol, however, this is of no interest for our usage scenarios. Being a passive RFID system, the chips do not need a power source. The reader is operated from a 4.5V power source and does not consume much energy, a further important attribute since the battery has to be carried by the user, too. In addition to the small antenna and reader, the chips themselves are small and cheap, which makes them

---

[15] See www.hitachi.co.jp/Prod/mu-chip/.

particularly useful for scenarios such as the supermarket.

To include the μ-chip system into the World Explorer prototype, two software components (in the sense described above) had to be written: one implementing the `BeaconSystem` component interface, the other the `InformationSystem` interface. The former – called simply `muchip` – encapsulates the tag reading protocol, while the latter – `muchipdatasource` – holds the XML information files. This external data source is necessary since the tags, unlike the motes, only possess an ID in a non-writable memory. The XML information file corresponding to this ID has to be retrieved from the external data source. Obviously, in a real deployment more than one external data sources would hold the information about the different products in a supermarket. Most likely is a scenario where a prefix of the standardized ID encodes the producer of the respective product, as proposed by the "Electronic Product Code" (ePC) initiative [53]. This prefix could then be used to derive the URL of the producer's information server (with a protocol similar to DNS, called "object naming service," ONS), and then retrieve from this server the information for the specific product, as identified by the rest of the ID.

Even without the RFID system included, the World Explorer was already cumbersome for the users, since it was occupying their only free hand (the other one usually being used for the cane or for holding the dog's leash). Since, after the addition of the μ-chip system, the user had to carry several components (the mobile device, the mote base station connected to one serial port, the RFID reader with battery connected to the other serial port, and the RFID antenna), a usability solution had to be found. Due to our idea of microscopic world disclosure, a first idea has been to mount the RFID antenna on the user's hand. By that, as soon as his hand would approach an object or a product, the antenna would discover it, and the user would hear the corresponding information, thus realizing the targeted explicit selection for such scenarios. Since, however, the antenna can only transmit its information to the reader via a cable, we have decided against this option due to usability issues.

In the end, we have realized a usability solution as depicted in Fig. 7.13. Hidden inside the backpack are the RFID reader and the mote base station, which do not need to be on the outside. On the left shoulder we placed the RFID antenna. Since it is thus placed close to the left ear (the same that we use for the audio output to the user via headset),

Figure 7.13: The usability solution for a system comprising both motes and RFID tags. On the left, the backpack containing the mote base station and the RFID reader. On its outside, the RFID antenna (left shoulder) and the mobile device (right shoulder). On the right, a typical usage of the system to scan a supermarket product for additional information.

from a user's perspective it seems that when a product is brought near the left ear (and is thus scanned by the antenna), it starts "whispering" to her. We have used the opportunity to sue a cover for the mobile device on the right shoulder of the backpack. Thus, when the user does not need to interact with objects, but only listens what they are, the mobile device can stay inside the cover, freeing the user's hands, and yet allowing her to learn about objects discovered by both the mote base station and the RFID antenna. When she decides to interact with an object, she can pull the device from the cover and interact in the usual way.

7.3.4 New Remote Control on the Cane

Despite the improvement in usability obtained by integrating the World Explorer's components in the backpack, one major usability problem remained. To interact with an object of the Chatty Environment (i.e., for getting more information about the object, or for triggering an action), the user had to use the only free hand for operating the mobile device. This, as our interviews had revealed, represents a major drawback for the blind and visually impaired. Furthermore, the joystick of the deployed PDA is rather imprecise, making it difficult for the user to perform the desired operation. Especially for the most frequent *select* operation, a more satisfactory solution was needed, one that would not be easily misinterpreted by the system.

To alleviate these usability issues, the last extension of the prototype has thus been the development of a remote control for the World Explorer that is included into the user's cane, and can be operated with the same hand already occupied by the cane. The interviews, however, had also told us that such a remote control must not interfere with the regular operation of the cane. One more pragmatical constraint being the fact that a cable leading from the cane to the backpack would obviously be cumbersome, it was clear that the so-called "smart cane" would have to transmit its results wirelessly to the World Explorer device.

Since this last system extension, although rather helpful for the users, and although it posed quite some implementation challenges, does not add novel conceptual issues, it will only be presented shortly. The first decision was which electronic device to use as remote control. Since both serial ports of the PDA running the World Explorer software were already used for the communication with the two tagging systems, and the system's WiFi connection was also already in use (for gathering information from external data sources), the only interface we could use was Bluetooth. Since, however, as stated at the beginning of section 7.3.2, one of the drawbacks of Berkeley motes is that they do not support Bluetooth, we could not use them for this task. The decision thus has been taken to use ETH's BTnode sensor node to transmit the commands to the World Explorer.[16]

As shown in Fig. 7.14, two buttons have been added on the prototypical "smart cane." One is a five-way navigation joystick, similar to the one on the mobile device, however, with much higher precision and also better usable with the thumb. The other is a single button on the other side of the cane's handle, which is a dedicated *select* button. The BTnode continuously polls the two buttons about once every millisecond. If any of them has been triggered, it sends the corresponding command to the World Explorer. To be able to do so, upon starting, the BTnode establishes an RFCOMM connection with the World Explorer on the

---

[16] See www.btnode.ethz.ch.

user's mobile device. The BTnode application has been written in C using the BTnode library.

Figure 7.14: The "smart cane," with the two command buttons (five-way navigation and dedicated *select* button) on the handle, an on-off button below, and the BTnode for Bluetooth command transmission further down. The batteries for the operation of the BTnode are hidden inside the cane's shaft. The picture's upper left corner shows a close-up of the two buttons for logical navigation on the handle.

On World Explorer side, a new `InputSystem` component has been written, called `stickinput`. During its initialization, the `stickinput` component opens the Bluetooth stack's virtual serial port for incoming connections (thus, when starting the BTnode, it will succeed to autonomously establish an RFCOMM connection).

7.4 Analyzing Scenario and Prototype - Insights

After having developed the prototype, we started analyzing the possible consequences of a real-life deployment of a system that discloses the environment to the blind and visually impaired. This analysis corresponds to the third step of the scenario-driven prototyping method, as summarized in Fig. 5.2. In order to do so, we presented the prototype to the same group of blind and visually impaired that we had interviewed before, and we further analyzed the technological consequences of scenario and prototype.

7.4.1 Main Features of a Chatty Environment

Among the most valuable insights generated by the project have been the preferences of potential blind and visually impaired users regarding the features of a system aiming at providing them more day-to-day independence. Their (for us at first surprising) preference for the world disclosure over the navigation property has already been presented.

In a second step, after the prototype has been completed, we again invited the blind and visually impaired from the first round of interviews to test our system. Eight of the nine original interviewees have responded to our invitation. They have been presented with two scenarios. In the first, users would move in an unknown environment (walk along corridors in an office building) and objects along the way would reveal themselves to them, such as the several offices and other rooms along the way, or objects like the coffee machine and the refrigerator in the kitchen. A total of 14 objects have been tagged by motes. The users had no specific task (like finding out where X is, or finding their way to Y), they could freely play with the system. The second scenario simulated a supermarket. The users were presented a shelf that contained five different sorts of (packed) spices and five different sorts of rice. These are items that feel the same and blind people are typically not able to distinguish among them while shopping. For some of the products, we had several packages with different expiration dates. Users had to find a specific item and – if unhappy with the expiration date – find another package that will be usable for a longer period of time.

Generally, the users liked the system. Especially the feature that it does not interfere with the normal usage of the cane has been highlighted by all participants. All of them could find the sort of rice they prefer, or spices with a longer expiration date without any noticeable problems. While walking through the environment, the users have extensively played with the system, getting information about the people in different offices, or about the history and features of the tagged objects.

When asked again a question from the first round of interviews, whether too much information would bother, the ratio of positive to negative answers almost inverted. While without holding the prototype in their hands, six out of nine said that too much information would ultimately become annoying (66%), now only two out of eight (25%)

would prefer to be able to adjust the number of objects that speak, i.e., the information density, depending on their actual needs. As one blind participant now put it, "there can never be too much information." While this change of mind might be due to the relative low density of tagged objects in our test, or to the novelty of the system, it might also point towards the large need for information of the blind and visually impaired, need that could be partly satisfied by such a system. As one female participant said, "although rather useless, it was really funny to learn when the refrigerator has been produced or which energy class it belongs to."

7.4.2 Issues when Designing an Interface for the Blind

While moving through the environment, information about objects in the user's neighborhood is gathered by his device. However, a series of questions arise regarding as to when to present this information, when to delete it, how to present it to the user and how to let the user logically navigate through an object or among several objects.

Push vs. Pull Model

At the beginning, the Chatty Environment only had a push model for presenting new objects to the user, i.e., as soon as the user's device senses a new object in the neighborhood, the user is informed of its presence. This seems the logical way to go – inform the visually impaired users about their surroundings without requiring them to take explicit actions. This paradigm also maps the way sighted people gather information about their surroundings – unobtrusively, from the corner of the eye, so to say. However, we soon realized that while this paradigm works well for large and relatively sparsely distributed objects, it fails for small objects with a high density like supermarket items. The user will certainly not want some hundreds or thousands of different products be announced to him as he enters the supermarket. Here, a pull model seems more appropriate – having the user explicitly choose an item, then supplying him the needed information. Technically, the two tagging systems work identical. As soon as a tag is discovered, the user is informed about its existence. However, we used the fact that the μ-chips have a reading distance of only a few centimeters. We placed the reading antenna (that is connected to the user's mobile device in his backpack) on one of the backpack's straps near the user's right ear.

Thus, from the user perspective, only when she brings such a product to her ear, it starts talking to her, the user being the one explicitly initiating the communication.

Logically Navigating through Objects

The user's mobile device is constantly polling the environment for Chatty Environment compliant object markers. After a new tag has been discovered, the information XML file is downloaded and the new object presented to the user. The information is organized as a tree that can be navigated by the user. The tree's root is the designation of the object. It can be freely chosen, but will usually be the object's textual description in a meaningful way for the user. Examples might be "Tram no. 9," "Ticket vending machine," or "Spice: Oregano."

When a newly discovered object is presented to the user, he will hear via earplug the text "New object" followed by the object's designation, e.g. "New object: ticket vending machine" (texts being converted to speech by the commercial text-to-speech engine installed on the mobile device). To understand how the user may logically navigate through an object or among several objects, lets take a look at the data model on the user's device (see Fig. 7.15). On the top level, there is a list of all objects active in the system. The information for an object is organized as a tree with the object's designation at its root. Trees can be of any height – the picture shows a tree of height two.

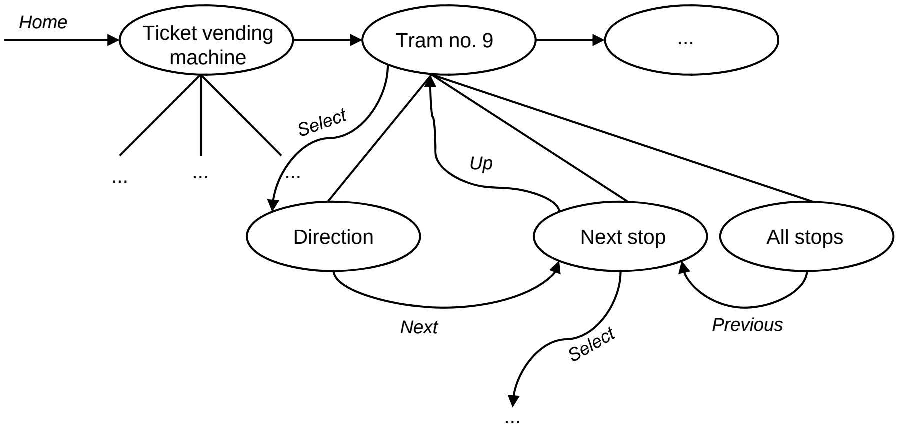

Figure 7.15: Example of the logical navigation through a list of objects at a tram stop.

Every time a new object is discovered, its designation is presented to the user. Then, the new object is added in front of the list of active objects. The user can always go down the branch of a tree if he presses the *select* button while the text of the father knot is being read, or for the subsequent two seconds. For example, if pushing the button while the system reads "tram no. 9" or shortly afterwards, the user will get one level deeper into the tree and will hear the next level of information for that object: "direction... [2 seconds pause]... next stop... [2 seconds pause]... all stops." Here again, he may choose to go one level further down by pressing the *select* button. These knots already being leafs in this example, the user will hear the desired information. If the user misses the opportunity to press the button, he may always navigate back on the same level by using the *back* button. He also has the opportunity of speeding up the process by pressing the *next* button. Going up one level in the tree can be done any moment by the *up* button. Also at any moment, the user might press the *home* button to get to the first object in the list (which is the last discovered). After pushing the home button, all active objects in the system will be read, starting from the newest one, until the user chooses one of them to "dive" into.

Deleting Objects from the System

This obviously raises the question of when to erase an object from the user's system. Since the mobile device is continuously polling the environment for markers, it always knows whether an object is still in range or not. While the object is still in range, no action will be taken; especially, the object's designation will not be presented to the user as being a 'new' one. A design decision that had to be taken, however, was how long to keep an object active after its signal has been lost by the mobile device. There are two reasons for not deleting objects at once. First, the user could move on the border of the object's signal radius, thus receiving and losing the signal repeatedly. If the object would be deleted at once after losing the signal, the user would be presented the same object again and again as a new one. Second, even after not being in range any more (but still not far away), the user might want to get some information about an object. On the other hand, keeping objects too long active might present the user an outdated world view and would contradict the paradigm of the Chatty Environment – presenting the user his neighborhoods. A trade-off has to be made here, which seemed rather difficult to be done globally. We learned that for motes a balanced default value would be 30 seconds.

For RFID tags, the default is set to 5 seconds. Every single object might overwrite this default though (by an attribute that can be set in the XML information file), since enforcing a global value for every possible application does not seem a viable option.

To realize this, every object in the system keeps a count of the time elapsed since communication with the corresponding tag has been last established. Every time the tag is found again, the counter is reset to zero. When the counter reaches its maximum value, the object gets deleted from the list on the top tree level and is not accessible by the user again (unless, of course, he moves again in its vicinity).

7.4.3 Context Awareness in a Highly Dynamic Context

While some of the data needed used by the World Explorer device is static (such as the expiration date or the ingredient list of products), other information relevant for the Worlds Explorer is changing, sometimes rapidly: the position of the user, the position of some of the tagged entities (such as trains or busses), or the state of some of the objects (such as a traffic light changing color). Since the system intends to disclose the environment to the user, it was thus one of the main requirements that the system presents him an up-to-date view of the environment. In some of the cases, for example when the user is being told the state of the traffic light, this requirement becomes imperative.

This feature of the World Explorer, that it needs to react to changing conditions in the environment, is by no means unique. Being able to sense their surroundings and often being mobile, a general attribute of ubiquitous computing applications is their capability to adapt to changes in the environment. This feature, one the most important concepts within ubiquitous computing, is called *context-awareness*. Few context-aware ubiquitous computing applications, though, have so hard constraints to cope with as the Chatty Environment does. This section, after introducing the concept, history, and definitions of context-awareness, presents thus the insights that we could generate from the project for the field of context-awareness.

Context Awareness

The term “context-awareness” has first been defined by Schilit et al.: “Context-aware software adapts itself to the location of use, the collection of nearby people, hosts, and accessible devices, as well as to changes to such things over time" [138]. In Schilit's definition, the location along with the implied neighborhood context are regarded as outstanding characteristics for defining the context of an application.

This definition, which focuses on the application itself, has later shifted towards a more user-centric definition, and has also became more general. The definition of context most widely accepted nowadays has been provided by Dey and Abowd: "Context is any information that can be used to characterize the situation of an entity. An entity is a person, place, or object that is considered relevant to the interaction between a user and an application, including the user and the application themselves" [42]. Preceding this general definition, several other contributions, such as Schilit's above-mentioned first definition [138], a previous contribution by Dey [43], and Pascoe's [122], have listed the most important aspects of context as:

- *User environment* – location, collection of people nearby, social situation.
- *Physical environment*, such as lightning conditions or acceleration of a vehicle.
- *Computing environment* – available processing power, communication protocols.

Such definitions, although not as general and thus sometimes too constraining, help on an operational level to identify the most common types of context.

Finally, Dey and Abowd further differentiate between *primary* and *secondary* types of context: "Location, identity, time, and activity are the primary context types for characterizing the situation of a particular entity" [42]. All secondary types of context can be derived from these fundamental ones: the email address, for example, from the identity of an entity, or the people nearby from its location.

We conclude this short presentation of context awareness with one more observation and one definition from Dey's and Abowd's work. The observation relates to implicit versus explicit context, and is relevant since most projects assume an implicit context. The authors, however, argue that explicit context is likewise important: "Our definition allows context to be either explicitly or implicitly indicated by the user. For example, whether a user's identity is detected implicitly using vision or explicitly via a login dialogue, the user's identity is still context" [42]. Finally, from their definition of context, the authors also define context-aware computing as "A system is context-aware if it uses context to provide relevant information and/or services to the user, whereas relevancy depends on the user's task" [42].

Context Awareness in the Chatty Environment

Context awareness obviously lies at the very core of the Chatty Environment concept and the World Explorer prototype. Providing the user with contextual information relevant to her task (which is, orient herself in an unknown environment) is the aim of the concept. The application continuously gathers environmental contextual data and presents it to the user, allowing the user to interact with the environment as well. Especially since the application has to cope with some relatively hard constraints, as will become clear shortly, throughout the development of the prototype we could thus gain a series of insights into the area of context-awareness. The two more important ones, a novel taxonomy of context-awareness systems and the concept of *dynamic information beacons*, are presented below.

Numerous context-aware ubiquitous computing applications have emerged over the last years. Assuming Dey's and Abowd's context definition [42], all context-awareness systems have to solve similar problems: find out which are the relevant entities that determine the user's and application's context, and gain information about these. Depending on how they solve this task, we differentiate between the following six types of systems.

Location with Static Information Lookup. The availability of the global positioning system (GPS) for civilian use by the mid-1990s triggered the emergence of a whole class of context-aware applications. These applications relate the user to his geographic context. They typically take as input the GPS coordinates sensed by an incorporated GPS receiver. As output, they graphically show the location of the user on a map, or navigate him to another point on the map. Examples for such systems are hiking maps for mountaineers, sailing and various other kind of electronic maps, and the well-known vehicle navigation systems.

From a technological perspective, all these applications rely on the same principle: the system's GPS receiver continuously determines the user's location, then this data is mapped to the local context (i.e., a point on the map). This correlation between location and points on the map is statically stored on the device, typically on a CD. Hightower and Boriello give a good overview of location-aware systems in [71].

The main problem with this approach is its lack of flexibility – if some of the information changes, all users have to replace their CDs with newer ones. If they don't, they will be confronted with outdated information. Therefore, this approach works fine for mountaineering maps, since mountain topology changes very slowly, over millennia. For vehicle navigation systems this is a feasible option, too, since street topology also changes rather slowly. But even for such systems, one is sometimes confronted with outdated information: glaciers that have changed their shape over a few years or road works not mentioned on the map.

Some of these systems provide a practical workaround: the user can update the data online from a central server before starting the journey. This workaround, however, only reduces the problem's quantity, not its quality. Real-world changes after the journey started will not be represented in the virtual world model the user carries around. Furthermore, another problem remains: The user has to pre-load data about the correct destination. When, for example, driving over the border, he has to possess the CDs for the destination countries as well.

Location with Online Information Lookup. The GUIDE project in Lancaster [38] was one of the first to put forward this new approach. The method has a key feature in common with static lookup systems: it first locates the users, then uses the location to retrieve information about the context. The difference between the two lies in the lookup method. In the first class of applications the lookup was locally done, the static data being preloaded to the system. Here, the system connects to a central server to retrieve the data. The user usually has no data preloaded on the system, although for pragmatic reasons there may be some caching involved.

The GUIDE project [38], for example, used the signal strength of WiFi access points scattered through the city center of Lancaster to determine the user's location. Being a tourist information system, it used the user's position to provide him relevant information about the neighborhood (historic buildings, restaurants, etc.). The lookup was done online at a tourist office server, over the same WiFi connection.

The advantage of such a system over static lookup systems is obvious: when information about an entity (e.g., its state) changes, the users must not update any locally stored data (since it does not exist); the information has to be refreshed only once at the server. The disadvantage over the static information lookup is that reliable connectivity to a communication infrastructure is needed.

While working fine for a series of applications, both location-based approaches, however, have two inherent drawbacks: Firstly, since the user's local context is derived from the user's location, these approaches cannot cope with moving objects (other than the user). They thus work fine for Lancaster castle or the highway topology of a country, but not for an approaching bus. Secondly, both paradigms are highly centralized. Entering new data about existing entities, or adding new entities of interest for the users, are tasks that can only be accomplished at the central instance – the one producer of navigational CDs, for the former method, or the administrator of the database for the latter.

Proximity with Static Information Lookup. Fundamentally different systems than those using location are ubiquitous computing systems that make use of *proximity* to learn about the user's neighborhoods and the implied local context. Hereby, we use the term "proximity" only in partial overlapping with the way it is used by Hightower and Boriello [72] ("Proximity measures nearness to a known set of points"). We rather mean proximity to other entities which are relevant for the user's context, without taking the detour over the location of those entities. Proximity systems place an electronic tagging device with a certain transmission range on relevant entities in the environment. When a user comes into the vicinity of such an entity, her mobile ubiquitous computing system receives the object's ID.

Proximity systems have two key advantages over location-based systems, and both are related to the implicitly existing neighborhood relation when the user's mobile device can receive the tags of an environment entity: Firstly, there is no need for any location infrastructure. Neither must the user be located to know what is around him, nor the objects themselves. Secondly, they can cope with object mobility in an intrinsic way, with no managing effort at all. In a location-based system, when a relevant entity changes its location, this fact has to be reflected at the lookup service, i.e., either locally on the user's device or on the lookup server. In proximity systems, when such an entity changes its position, this fact will reflect immediately and inherently into the system: users will not receive any transmission at the old location, but in the vicinity of the new location only.

The same differentiation between static lookup (on a data source carried by the user), and online lookup can be made here, too. One example for a proximity system with static lookup is another world disclosure system for the blind, “Navigational Assistance for the Visually Impaired” (NAVI) [123]. In NAVI, the blind user’s mobile device is equipped with an RFID reader and a CD containing the data about objects that will be discovered in the user’s environment. Despite the two advantages mentioned above (no need for location system and suitability for moving objects), such system suffers from the same major drawback as location-based systems with static lookup: its lack of information flexibility.

Proximity with Online Information Lookup. To counter this last drawback, proximity systems with online information lookup, such as the RFID part of the World Explorer, gather their data from online sources, which can be updated more easily (on server side). More importantly, such systems are decentralized: anyone can spread tags in the environment, and provide the corresponding database with data. Note that this is an advantage that location-based systems with online lookup did not offer over their static lookup counterparts: There, since the (online) lookup system is correlated to one database only (the database of the one location system), the system is still centralized.

Static Information Beacons. Substantially different than the paradigms presented above are information beacons. These more powerful tagging devices directly store the information about the object they tag.

One classic example of such applications is IRREAL [10], developed at the University of Saarbrücken. IRREAL is an indoor navigation and information system. Infrared transmitters located throughout the building continuously transmit locally relevant contextual information, such as a map of that part of the building. Only building information have been sent in the IRREAL project, it is however conceivable and conceptually identic that other objects could broadcast information about themselves, too. In the same way, should these other objects be mobile, their movements would naturally be represented in the system, as stated above. Thus, the word “static” does not refer to “beacons” but to "information." Relevant for this distinction is not whether the entity the beacon is placed on is mobile or not, but whether the information sent by the beacon is always the same or not.

Static information beacons have several advantages in common with proximity-based systems: They do not need any location system, and they inherently cope with mobility. Additionally, they combine two key advantages of the two flavors of proximity-based systems: as the proximity-based systems with online lookup, they are decentralized. Anyone can deploy a beacon in the real world, which already contains the information about the object it tags. Furthermore, they are as independent from a communication infrastructure as the proximity-based systems with static lookup. Since all relevant data is already stored on them, the user's mobile device does not need to connect to remote databases through protocols such as WiFi, GSM, or UMTS.

Dynamic Information Beacons. The main disadvantage of static information beacons, though, is the static nature of the information stored on them. As with the other paradigms using static data, updating the information to reflect changes in the real world is costly, will happen with a certain pace, and thus parts of the system's worldview will always be outdated. In other words, such systems are ill-suited to "bridge the gap between real world and virtual world" [132]. Arguably, updating such beacons is even harder than CDs with static data about an environment, since they already have been deployed in the real world.

In the Chatty Environment, we have foreseen the need to have dynamic, up-to-date information about the environment being transmitted to the user, and have thus implemented the concept of *dynamic information beacons.* The beacons we have used can be augmented with sensors and can include the data gathered by these into the XML file they are continuously sending in the environment. Furthermore, we have foreseen that some of the information thus sent might be urgent, such as the change of a traffic light from green to red. Even if the blind user's device has already presented the presence of a (green) traffic light to its user, in the moment the color changes to red, the mote on the light pushes a new, high priority XML file on the user's device. Upon arrival of a high-priority message, the device interrupts its current action, independent of its nature, and presents the new message to the user.

Context Awareness Criticism

As a final observation related to context-awareness, the concept has recently increasingly been questioned. Rogers, for example, who questioned Weiser's vision of calm technology [165] in a recent controversially disputed article [130], argues that ubiquitous computing research should engage the user more, writes about context: "Ethnographic studies of how people manage their lives – ranging from those suffering from Alzheimer's disease to high-powered professionals – have revealed that the specifics of the context surrounding people's day-to-day living are much more subtle, fluid and idiosyncratic than theories of context have led us to believe. This makes it difficult, if not impossible, to try to implement context in any practical sense and from which to make sensible predictions about what someone is feeling, wanting or needing at a given moment" [130].

By highlighting explicit context (which engages the user) into their broad definition of context as early as 2000, Dey and Abowd [42] have already and well ahead of time precluded this criticism. We have also proposed with the Chatty Environment a "smart" application that does not tap into the trap of artificial intelligence. Our system informs the user about pasttimes, for example, but does not try to take any decisions as of his (obviously complex and impossible to model in a computer system) preferences.

7.5 Related Work

There has been no significant effort to address blind navigation other than using the cane and a guide dog until the early 1990's. Then, satellite positioning services such as the Global Positioning System (GPS), and new generations of computers with sufficient computational power opened new perspectives to help blind people move through the environment more safely and easily. Several research projects on GPS-based guidance systems with speech output thus started. One of the first prototypes, called "Personal Guidance System" [103], has been developed by the group around Prof. Jack M. Loomis from the University of California, Santa Barbara. The system runs on a laptop computer carried in a pack slung over a shoulder and weights a couple of pounds. Drishti [68] is a similar system developed at the University of Florida. It computes optimized routes based on user preferences. The computer, the GPS-receiver, and an electronic compass are carried in a backpack. Michael May's Senderogroup developed "GPS Software v2,"[17] as an extension to their PDAs for blind people called "BrailleNote" and "VoiceNote." An external GPS receiver connected to the PDA sends position information to the software that is able to provide relative distances and directions to points of interest and to create routes for pedestrians, in a similar manner to car-based navigation systems.

One of the first and still rare indoor orientation systems for blind people is "PERSONA" (personal, electronically-recorded, speech-output navigation aid), developed at the Jerusalem College of Technology under Yehuda Sonnenblick in 1998 [145]. It consists of several purpose-built transmitter units that continuously send a unique location code using infrared beams, and of portable receiver units with an integrated speech module. A more recent indoor navigation system for the visually impaired, using the signal strength of already installed off-the-shelf WiFi access points, has been developed at the University of Stuttgart, Germany [80].

Systems with world disclosure features, on the other hand, are rare. An early system to implement such functionality is Talking Signs,[18] which started back in 1993. Infrared beacons mark the environment's objects, the user's mobile device has an infrared receiver and a small computing platform. The used infrared technique has, however, several drawbacks. The user has to actively scan the environment, pointing the device to all possible directions until an answer comes, which is obviously a large drawback and often unpracticable for the blind. Furthermore, while scanning the environment, one of the user's hands is occupied. Most important, however, since only short audio messages are transmitted, the user cannot navigate through the object to gain more information about it.

The more recent system "Navigational Assistance for the Visually Impaired" (NAVI) [123] also uses a similar approach. The user's portable device combines a CD player with a mobile RFID tag reader. The tags mark objects in the environment and trigger the corresponding track on the CD. The approach is similar to ours – the device explains the surroundings to the user by reading messages such as "front of Rush Rhees library" [123]. However, NAVI has the drawbacks of proximity systems with static information lookup. First, the system does not

---

[17] www.senderogroup.com/shopgps.htm.

[18] www.talkingsigns.com.

scale. The user has to know a priori where she is heading to and insert the correct CD (which she must have obtained in advance). When the number of tagged entities in a given environment varies, the CD must be updated, too, and supplied to all users. Second, using passive RFID tags only constraints the perception of all tagged objects to a radius of about one meter (depending on the used tags and reader). It is often advantageous, however, to be able to define virtual auras with different extensions for different classes of objects (since, most likely, a railway station should have a larger aura than a package of baked beans). Third, NAVI's approach does not allow objects to change their state, since the information is stored statically on the CD.

There are further some prototypes for the aid of the blind or visually impaired that target world disclosure for a specific area only. One such system is the "closet buddy" [133] that uses RFID tags sewed into clothes to retrieve information about their colors and gives advices to the blind on how to match them. For this specific application, the system thus exhibits a world disclosure property similar to our system. A better known and more flexible system, conceived for helping blind in public transportation scenarios, is Ubibus [9].

7.6 Summary

In this chapter, we presented a second project in which we have applied the method of scenario-driven prototyping introduced in chapter 5. In a first step of the project, we identified within the multidisciplinary dialogue the large potential that ubiquitous computing possesses in the field of assistive technologies. Taking advantage of first insights into the everyday problems of blind and visually impaired given to us by a blind colleague, we wrote a scenario envisioning how these could be alleviated by the use of the novel technologies. Thus, as opposed to the technology-driven Smart Tachograph scenario, the scenario of this chapter is a problem-oriented scenario.

The scenario envisions the enhancement of today's visual information by other means of information that can be experienced by the visually impaired, mainly by spoken information. Through the electronic enhancement of objects with markers that can hold information about the object they tag, a mobile system carried by the user could pick up this information when coming in their vicinity. It would then be presented as spoken information to the blind or visually impaired user, who would thus more easily find his way in unknown environments. Due to this feature, we called this vision the "Chatty Environment." Furthermore, using this information as basis, the system is also able to guide the user to points of interest. If, for example, the system discovers a poster announcing a symphonic concert, the user is not only able to gather more information about the place, hour, and conductor, but also to ask the system to navigate him to the concert hall or to the nearest ticket selling facility.

While developing the scenario, we were almost convinced that this second feature, *navigation*, would be the one of real interest to the blind users, while the former, *environment disclosure*, would mainly be a prerequisite for attaining it. When conducting a series of interviews with potential users, however, the majority of blind and visually impaired told us that the environment disclosure property could actually be the more important one, the reason being their great need for information. This need of information is not only of psychological nature, it also has practical purposes. Except the obvious everyday problems of blind that any sighted researcher can quickly understand, the potential users pointed us towards less obvious problems they encounter in their daily routine, such as the difficulties when shopping for food packages that all feel the same, or in catching a bus whose doors only open when a button is pushed.

Following the paradigm of scenario-driven prototyping, we thus developed a first prototype, focusing exclusively on the world disclosure property of the Chatty Environment. The prototype, "World Explorer," which has been thoroughly presented, tries to match the way sighted persons perceive the world. To this respect, we used as tagging devices for real-world objects the Berkeley motes, small computation and communication devices originally designed for sensor networks applications. The motes have several advantages: they have writable storage, so that they can store information files about the objects they tag, they have a large transmission range of about one hundred meters, and, most importantly, this range can be fine-granularly adjusted by adjusting the transmission power. Through this last feature we could obtain our goal of mapping the way sighted people perceive their environments: large or important objects will send the information about their existence and their attributes at a greater distance, while smaller objects will have a smaller transmission range.

Since this push-model (the object is pushed to the user's attention) is well-suited for relatively large and sparsely present objects (the "macro-world"), but less adequate for dense populations of small objects such as supermarket products (the "micro-world"), we have then extended the prototype as to allow passive RFID tags as well. Although technologically the interaction is identical, through their very small range of a few centimeters, the RFID tags allow a different interaction paradigm from the user's perspective – a pull model.

Through the tight correlation of scenario and prototype development required by the scenario-driven paradigm, as well as the multidisciplinary dialogue and stakeholder interviews needed, we could thus generate a broad set of novel insights. They reside more on the technological side, but are heavily influenced by the societal findings and preferences of the involved potential users. The first insight has been the preference expressed by the users for world disclosure over navigation. Coming as a surprise to us sighted researchers, this fact probably also explains why there are quite a number of assistive technology projects that aim at providing navigational aids for the blind, but only a few with environment disclosure at their core.

We could further identify some novel HCI aspects during the design of a speech-and-choice interface, which are likely to have relevance beyond our project for any assistive system for the blind. We have recognized the already mentioned differences from the user's perspective between a push and a pull model, and when each of these is more appropriate. We have proposed a novel solution for the logical navigation of an environment, which has been received favorably by both the potential users and the academic HCI community. Finally, we have addressed the tradeoff between letting out-of-range objects be active for a longer time, risking thus to present the user an outdated view of the environment, versus deleting them at once, with the risk of "continuously reinventing the wheel."

A further contribution came from the hard requirements posed to the context awareness of the system. The prototype needs to be able to cope with dynamic environments, to which newly tagged objects will continuously be added. Many of the context-relevant objects will be mobile, thus a simple location-based approach will not suffice. Furthermore, some of these entities will change their state, possibly quite often, and these changes will be of importance for the users. Finally, some of the contextual information will have high priority, and will have to be conveyed to the user at once. We approached the problem by first proposing a novel taxonomy of context-aware systems, which differentiates alongside four axes: local vs. online information lookup, location-based systems vs. proximity-based systems, tags (that only hold an ID, but not the information itself) vs. information beacons, and static vs. dynamic information (i.e., sensing the environment). Our prototype occupies two formerly sparsely populated spots inside this taxonomy: the RFID-based part of the Chatty Environment is one of few proximity-based online lookup services, while the Chatty Environment part based on motes has been the first proposal of dynamic information beacons.

8 Conclusions

In this thesis, we have shown that the expected proliferation of ubiquitous computing technology is likely to have a profound societal impact and to pose an important number of technical challenges. After an in-depth analysis of the technological advances that contribute to this proliferation, we have presented the economic reasons that speak in its favor as well. This brings a wide panoply of societal opportunities and risks, which has been thoroughly exemplified.

Because of the expected societal importance, the field has started to enter into the public consciousness and to be confronted with numerous critical questions. The assessment of ubiquitous computing implications thus becomes not only important for providing answers to these critical questions, but also for identifying the technological challenges that can lead to the alleviation of the most urgent societal concerns. However, we have further argued that neither the existing technology assessment methods could provide satisfying results so far, nor most prototypes developed by ubiquitous computing researchers. The former suffer from the difficulty in coping with the broadness and the high technological dynamism of the field, while the latter typically lack the broad scope needed for a general societal view. To alleviate these problems, this thesis makes three contributions:

- *Scenario-driven prototyping*, a novel method for the assessment of societal and technological implications of ubiquitous computing, which tightly brings together scenario analysis and prototype development within a four-step multidisciplinary process;
- a first case study, in which the method has been applied, the *Smart Tachograph*, a project using ubiquitous computing technologies for the individual assessment and distribution of traffic costs; and
- a second case study, the *Chatty Environment*, a ubiquitous computing project which aims at providing more everyday independence to the blind and visually impaired by auditorily revealing them their environment.

This final chapter briefly summarizes the three contributions, shows the limitations of our method and points to directions of further research.

8.1 Contributions

In this section, we summarize the three main contributions of our work. These have also been published, most notably in [16, 17, 25, 26, 28, 29, 30, 33, 34, 35, 36].

8.1.1 Scenario-Driven Prototyping Method

To avoid the development of prototypes with little societal relevance by engineers, or the technologically unsophisticated studying of implications by social scientists, we have proposed in this thesis a close multidisciplinary coupling of the two, inside a method that we call *scenario-driven prototyping*. This novel paradigm distillates and structures the experiences gathered and the actual developments accomplished while managing the three-year interdisciplinary project "Living in a Smart Environment – Implications of Ubiquitous Computing."

Scenario-driven prototyping proposes the development of both technology assessment scenarios and of prototypes emphasizing the main aspects of the scenarios, in a process encompassing four steps: Firstly, in a multidisciplinary dialogue, the engineers present the possibilities of the technology and the likely developments in the near future. Based on this presentation and on the societal issues and future conflict potentials identified by social scientists, they together choose relevant application domains and particular applications inside these domains. The result of this first step are future scenarios. In the second step, one or more prototypes are developed for each scenario, underlining the scenario's main aspects. Thirdly, new technological and societal issues are identified by analyzing scenarios and prototypes. The analysis should again involve a multidisciplinary approach, and, if appropriate, stakeholders should be interviewed. The so-discovered issues influence back in the fourth step both scenarios and subsequently prototypes. This feedback loop can be repeated several times.

We have claimed that, aside of unifying specific advantages of scenario analysis and prototype development, our method also provides some advantages that can only arise from the mutual insemination between technological development and societal analysis. The two most important synergetic advantages are: Firstly, being derived from scenarios written with the broad societal picture in mind, the developed prototypes are also likely to emphasize a societally relevant opportunity or risk. The technological issues raised will thus either contribute to encouraging the former or discouraging the latter. Secondly, the scenarios will have the "technology check" needed to avoid the risk of being technologically naive and analyze opportunities and risks of impossible applications.

8.1.2 The Smart Tachograph

In the very beginning of one of the two presented case studies, the Smart Tachograph, the value of the scenario-driven prototyping method became clear. Although we had found in a first multidisciplinary dialogue an application domain – the expected proliferation of dynamic and individual prices – with likely important societal consequences and related technological challenges (context awareness of products, complexity management, novel interfaces), in further dialogues with economists and technology philosophers it became clear that due to acceptance issues we were probably heading towards a dead-end road.

Instead of developing a prototype that would have resulted in an, at least partially, irrelevant societal analysis and a potentially inadequate (since never needed) solution of technological issues, we have sought a more relevant domain to apply the discovered concept of dynamic and individual pricing. The resulting paradigm of a Smart Tachograph, an infrastructure for the individual and behavior-dependent accounting of traffic costs, has not only proven to be a great deal more realistic from an economic and societal viewpoint, but it has also made several novel contributions, both technological and societal (among the latter, especially economic).

The few years that have passed since the project's ending have witnessed a continuous increase in the interest shown by the insurance industry, the politics, and the public discourse in two subtopics of the Smart Tachograph: behavior-dependent vehicle insurances and road pricing. In this context, two of the main contributions of our project have gained an increased importance. For one, we have shown that a generic system used for both the accounting of individualized insurance premiums, and the payment of a context-dependent road tax (a system that can further be extended to any other type of traffic costs), is not only technologically feasible, but also leads to a simplified overall solution. Secondly, we could prototypically document that such systems, with a potential high societal value but an apparent inherent price in terms of privacy intrusion, can be implemented with a high degree of respect for the driver's personal privacy, contrary assertions from the insurance industry notwithstanding.

8.1.3 The Chatty Environment

The other case study we presented is the Chatty Environment, a paradigm which makes use of the ubiquitous computing technology spread in the environment to enable blind and visually impaired people to lead a more independent life. Again, the tight coupling of scenario writing, stakeholder survey, and prototype development, has started to show its value at an early stage: We could learn through interviews with blind and visually impaired persons as potential users that from the foreseen two functionalities of the system, not the one we had expected to be more important, navigation through unknown environments, but the other, environment disclosure, is likely to be more valuable to the users.

We have thus focused our implementation efforts in this area and could, besides the interest of potential users, gain valuable technological insights into the domain we were facing. Since a system constantly unrevealing the environment to the blind user has to cope with rapid changes in this local context, we have proposed a taxonomy for context-aware applications and made an initial contribution into one of these application classes – context-aware systems based on dynamic information beacons. While this concept has emerged in the area of assistive technologies for the blind, its findings can influence numerous other fields, such as electronic tourist guides, as well. We have further conceptualized the difference between macro-world and micro-world disclosure and the two different interaction paradigms needed by these: push and pull. Finally, we could propose a novel interaction paradigm for the blind – listen and select – and especially define a consistent way of logically navigating through information about the surrounding objects using this interface.

8.2 Limitations

With the two case studies presented, we have only provided *qualitative* proof for our claims, but no *quantitative* proof. In order to have done the latter, we would have needed to prove the superiority of our method as compared to others. This would have required to analyze future implications of ubiquitous computing technology and identify relevant technological challenges with our method as well as with others, and to subdue the so-gathered findings to the test of time (i.e., measure at a future point, when the implications will have fully developed, the results obtain through each method). This was not practicable for several reasons, out of which the most important are the following two. The test of time would need a time horizon larger than the one of our thesis. Furthermore, quantitatively measuring the results would need a quantification of the extent to which the findings from each method have come true, how important these findings are, and how important other, non-disclosed implications or challenges (for each method) have proven to be. All these criteria seem, however, difficult to quantify. As Hilty writes, referring to one of these issues, the quantitative measurement of pervasive computing risks, “We must conclude that quantitative risk assessment is not a suitable approach for the type of study discussed here. Instead, it seems necessary to use qualitative criteria to evaluate potential risks” [77]. Nevertheless, some of the insights generated by the two projects conducted according to the method have already, as presented above, proven to be relevant societal issues, or technological challenges, respectively.

As section 5.3 has shown in more detail, further limitations of the scenario-driven prototyping method include: the longer time needed to accomplish the analysis (due to the request for both scenarios and prototypes, and the necessary feedback loop), and the technologically limited peek into the future (since the technologies depicted have to be at least prototypically available due to the requested grounding into prototypes). Furthermore, since it has been tailored to the needs of ubiquitous computing technology assessment, it is yet unclear how well our method could successfully be applied outside this field.

8.3 Future Work

Our experiences with conceiving the method, applying it in two large projects, and implementing comprehensive prototypes for those projects, suggest several desirable possibilities for further work.

8.3.1 Scenario-Driven Prototyping Method

The scenario-driven prototyping method, with the four steps it requires, stresses the importance of tight connection between scenario development and prototype development, and puts forward the best way to do so, according to our experience of multidisciplinary ubiquitous computing technology assessment. However, the thesis is a computer-science thesis from a ubiquitous computing professional, and it does not make any recommendations as to how the scenario analysis has to be carried out.

Therefore, regarding the scenario-driven prototyping method itself, the most meaningful complement would be the inclusion of a recommendation as to how the multidisciplinary analysis leading to the scenario has to be performed. Most promising candidates hereby would be: a formalization for the process used for the choice of fields that will be looked upon, as well as the ubiquitous computing risk filter introduced by Hilty [75], which would formalize the choice of potential opportunities and risks inside these fields.

8.3.2 Smart Tachograph

While our current Smart Tachograph prototype implementation already supports all software concepts for the realization of client-side personalization, it lacks a mechanism for the tamper-proof distribution of software, as a real deployment of such a system would require. Through the inclusion of a trusted computing platform, the black box would effectively ensure the correct deployment of the insurer's billing code. To also counter man-in-the-middle or replay attacks from the user, the system should further include a public-key-infrastructure, with secret/public key pairs for all black boxes, or customer-specific keys given by every accounting authority to its users. To increase the system's degree of realism even more, interesting further research can be done to counter possible sensor tampering attacks. To this respect, not only hardware measures could prove efficient, but also computer science concepts, as we have described in section 6.4.2. The inclusion of data mining algorithms inside the black box, or of data perturbation techniques, while at the same keeping the high level of privacy granted by the current version of the prototype, represent interesting research avenues to be pursued in the future.

8.3.3 Chatty Environment

The World Explorer prototype comprehensively covers the world disclosure property of the Chatty Environment, including two supported tagging technologies, which complement each other in covering what we have called "macroscopic world" and "microscopic world," together with the different push and pull disclosure paradigms needed by these. It has further contributed to several insights about the logical navigation through real-world objects by blind or visually impaired users. The prototype does not, however, implement any of the navigational functionality foreseen for the Chatty Environment. Implementing this functionality would profit not only the technological research about interface design for the blind and visually impaired, and context awareness, but also the stakeholder discussion, which, as our experience has shown, brings a large amount of additional insights when the potential users can test the system's functionality. It would also make use of the interesting insights already gathered from blind and visually impaired interviewees regarding their priorities for the chosen path, which partly differ substantially from the priorities of sighted people. Including navigation would also enhance the system with the synergetic advantages of the complementing world disclosure and navigation properties. Numerous research questions seem to raise from this joining of the two functionalities, including speech-and-choice interface questions (how to include the navigation functionality in the exiting interface), ontology questions (how can the user choose targets to be guided to, which are not yet active in his system), route compilation algorithms (taking into account the priorities of the blind), and context awareness (compiling these routes accounting for several real-world attributes, such as the position of moving obstacles).

8.4 Final Word

In the preface of his 1944 book "What is life?," which has reportedly influenced one of the discoverers of the double helix to change his biological field of study towards genetics [57], the physicist Erwin Schrödinger[1] writes: "We have inherited from our forefathers the keen longing for unified, all-embracing knowledge. The very name given to our highest institutions of learning reminds us, that from antiquity and throughout many centuries the *universal* aspect has been the only one to be given full credit. But the spread, both in width and depth, of the multifarious branches of knowledge during the last hundred odd years has confronted us with a queer dilemma. We feel clearly that we are only now beginning to acquire reliable material for welding together the sum total of all that is known into a whole; but, on the other hand, it has become next to impossible for a single mind fully to command more than a small specialized portion of it. I can see no other escape from this dilemma (lest our true aim be lost for ever) than that some of us should venture to embark on a synthesis of facts and theories, albeit with second-hand and incomplete knowledge of some of them – and at the risk of making fools of themselves" [139].

Obviously, we do not claim to have made through this thesis any contribution that could be compared in the least to Schrödinger's revolutionary work. Should we, however, have sometimes put forward "second-hand or incomplete" knowledge about the societal issues we have ventured into from the known grounds of computer science, may the reader generously overlook it in Schrödinger's universality spirit.

---

[1] Schrödinger had already received a Nobel prize in physics when expanding his interests towards biology.

Bibliography

[1] Nabil R. Adam and John C. Worthmann. Security-control methods for statistical databases: A comparative study. 21(4):515–556, 1989.

[2] Natascha Adamowsky. Smarte Götter und magische Maschinen – zur Virulenz vormoderner Argumentationsmuster in Ubiquitous-Computing-Visionen. In Friedemann Mattern, editor, *Total Vernetzt: Szenarien einer informatisierten Welt*. Springer, Berlin Heidelberg New York, 2003.

[3] Jari Ahola. Ambient Intelligence. *ERCIM News*, (47), October 2001.

[4] George Akerlof. The market for lemons: Qualitative uncertainty and the market mechanism. *The Quarterly Journal of Economics*, 84(3):488–500, 1970.

[5] Petteri Alahuhta, Paul De Hert, Sabine Delaitre, Michael Friedewald, Serge Gutwirth, Ralf Lindner, Ioannis Maghiros, Anna Moscibroda, Yves Punie, Wim Schreurs, Michiel Verlinden, Elena Vildjiounaite, and David Wright. Dark scenarios in ambient intelligence: Highlighting risks and vulnerabilities. Technical report, The European Union IST (Information Society Technologies) Advisory Group, 2006.

[6] S. J. Allin, A. Bharucha, J. Zimmerman, D. Wilson, M. J. Roberson, S. Stevens, H. Wactlar, and C.G. Atkeson. Toward the automatic assessment of behavioral disturbances of dementia. UbiHealth 2003: The 2nd International Workshop on Ubiquitous Computing for Pervasive Healthcare Applications, October 2003.

[7] Agustin A. Araya. Questioning ubiquitous computing. In *Proceedings of the 1995 ACM 23rd Annual Conference on Computer Science*. ACM Press, 1995.

[8] William A. Arbaugh, David J. Farber, and Jonathan M. Smith. A secure and reliable bootstrap architecture. In *Proceedings of the IEEE Symposium on Security and Privacy*, pages 65–71, May 1997.

[9] Michel Banâtre, Paul Couderc, Julien Pauty, and Mathieu Becus. Ubibus: Ubiquitous computing to help blind people in public transport. In *Proceedings of MobileHCI 2004*, pages 310–314, Glasgow, Scotland, September 2004.

[10] Jörg Baus, Antonio Krüger, and Wolfgang Wahlster. A resource-adaptive mobile navigation system. In *Proceedings of IUI'2002: International Conference on Intelligent User Interfaces*, San Francisco, California, USA, January 2002. ACM Press.

[11] Michael Beigl, Tobias Zimmer, Albert Krohn, Christian Decker, and Philip Robinson. Creating ad-hoc pervasive computing environments. Video contribution at Pervasive 2004, April 2004.

[12] Wendell Bell. *Foundations of Futures Studies – Vol.1: History, Purposes, Knowledge.* Transaction Publishers, 2003.

[13] Claudia Binder and Svetlana Domnitcheva. When waste becomes intelligent: Assessing the environmental impact of microchip tagging. In *Proceedings of SETAC/ISIE/Swiss Discussion Forum 2003*, Lausanne, Switzerland, December 2003.

[14] Eddie Blass. Researching the future: method or madness? *Futures*, 35(10):1041–1054, 2003.

[15] Jürgen Bohn. The smart jigsaw puzzle assistant: Using RFID technology for building augmented real-world games. Workshop on Gaming Applications in Pervasive Computing Environments at Pervasive 2004, Vienna, Austria, April 2004.

[16] Jürgen Bohn, Vlad Coroamă, Marc Langheinrich, Friedemann Mattern, and Michael Rohs. Disappearing computers everywhere – living in a world of smart everyday objects. In *Proc. of New Media, Technology and Everyday Life in Europe Conference*, London, UK, April 2003.

[17] Jürgen Bohn, Vlad Coroamă, Marc Langheinrich, Friedemann Mattern, and Michael Rohs. Living in a world of smart everyday

objects – social, economic, and ethical implications. *Journal of Human and Ecological Risk Assessment*, 10(5):763–786, October 2004.

[18] Jürgen Bohn and Harald Vogt. Robust probabilistic positioning based on high-level sensor-fusion and map knowledge. Technical Report 421, ETH Zurich, April 2003.

[19] Price Discrimination (June 30, 2007, 15:59 UTC). In *Encyclopaedia Britannica Online*. Encyclopaedia Britannica, www.britannica.com/eb/article-9061344/price-discrimination, 2007.

[20] Ryland C. Bryant, Eric J. Seibel, Cameron M. Lee, and Konrad E. Schroder. Low-cost wearable low-vision aid using a handmade retinal light-scanning microdisplay. *Journal of Society for Information Display*, 12(4):397–404, 2004.

[21] Lillian N. Cassel and Ursula Wolz. Client side personalization. In *Proceedings of DELOS Workshop: Personalisation and Recommender Systems in Digital Libraries*, 2001.

[22] Angela Chang, Ben Resner, Brad Koerner, XingChen Wang, and Hiroshi Ishii. Lumitouch: An emotional communication device. *CHI2001 Extended Abstracts*, pages 313–314, 2001.

[23] Hyemin Chung, Chia-Hsun Jackie Lee, and Ted Selker. Lover's cups: Drinking interfaces as new communication channels. *CHI2006 Extended Abstracts*, pages 375–380, 2006.

[24] Sunny Consolvo, Peter Roessler, and Brett E. Shelton. The carenet display: Lessons learned from an in home evaluation of an ambient display. In Nigel Davies, Elizabeth Mynatt, and Itiro Siio, editors, *Proceedings of the 6th International Conference on Ubiquitous Computing: UbiComp 2004*, pages 1–17. Springer, 2004.

[25] Vlad Coroamă. The chatty environment – a world explorer for the visually impaired. In *Adjunct Proceedings of the 5th Conference on Ubiquitous Computing (UbiComp 2003)*, Seattle, WA, USA, October 2003.

[26] Vlad Coroamă. Experiences from the design of a ubiquitous computing system for the blind. In *Adjunct Proceedings of the CHI*

2006 Conference on Human Factors in Computing Systems, Montréal, Canada, April 2006.

[27] Vlad Coroamă. Pervasive Computing im Alltag – Realistische Zukunftsanwendungen zur Untersuchung von Chancen und Risiken autonomer, intelligenter Objekte. *digma - Zeitschrift für Datenrecht und Informationssicherheit*, 6(03):106–109, September 2006.

[28] Vlad Coroamă. The smart tachograph – individual accounting of traffic costs and its implications. In *Proceedings of Pervasive 2006*, Dublin, Ireland, May 2006.

[29] Vlad Coroamă, Jürgen Bohn, and Friedemann Mattern. Living in a smart environment – implications for the coming ubiquitous information society. In *Proceedings of the International Conference on Systems, Man and Cybernetics 2004 (IEEE SMC 2004)*, pages 5633–5638, The Hague, The Netherlands, October 2004.

[30] Vlad Coroamă, Jörg Hähner, Matthias Handy, Patricia Rudolph-Kuhn, Carsten Magerkurth, Jürgen Müller, Moritz Strasser, and Tobias Zimmer. Leben in einer smarten Umgebung: Ubiquitous-Computing-Szenarien und -Auswirkungen. Technical Report 431, ETH Zurich, December 2003.

[31] Vlad Coroamă and Matthias Handy. Wohin verschwindet der Computer? Ein kontroverser E-Mail-Wechsel. In Friedemann Mattern, editor, *Die Informatisierung des Alltags – Leben in smarten Umgebungen*, pages 329–350. Springer, Berlin Heidelberg New York, 2007.

[32] Vlad Coroamă and Norbert Höckl. Pervasive insurance markets and their consequences. *First International Workshop on Sustainable Pervasive Computing at Pervasive 2004*, April 2004.

[33] Vlad Coroamă, Tarik Kapic, and Felix Röthenbacher. Improving the reality perception of visually impaired through pervasive computing. In Alois Ferscha, Horst Hoertner, and Gabriele Kotsis, editors, *Advances in Pervasive Computing*, pages 369–376. Austrian Computer Society (OCG), Vienna, Austria, April 2004.

[34] Vlad Coroamă and Marc Langheinrich. The smart tachograph.

Video submission abstract. Adjunct Proceedings of UbiComp 2005, September 2005.

[35] Vlad Coroamă and Marc Langheinrich. Personalized vehicle insurance rates – a case for client-side personalization in ubiquitous computing. In *Workshop on Privacy-Enhanced Personalization at CHI 2006*, Montréal, Canada, April 2006.

[36] Vlad Coroamă and Felix Röthenbacher. The chatty environment – providing everyday independence to the visually impaired. Workshop on Ubiquitous Computing for Pervasive Healthcare Applications at Ubicomp 2003, October 2003.

[37] Nigel Davies, Keith Cheverst, Keith Mitchell, and Adrian Friday. 'caches in the air': Disseminating tourist information in the guide system. In *Proceedings of the 2nd IEEE Workshop on Mobile Computing Systems and Applications (WMCSA'99)*, New Orleans, Louisiana, USA, February 1999. ACM Press.

[38] Nigel Davies, Keitqh Cheverst, Keith Mitchell, and Adrian Friday. 'caches in the air': Disseminating tourist information in the guide system. In *Proceedings of the 2nd IEEE Workshop on Mobile Computing Systems and Applications (WMCSA '99)*, New Orleans, Louisiana, USA, February 1999. ACM Press.

[39] Nigel Davies and Hans-Werner Gellersen. Beyond prototypes: Challenges in deploying ubiquitous systems. *IEEE Pervasive Computing*, 1(1):26–35, 2002.

[40] Robbie E. Davis-Floyd. Storying corporate futures: The shell scenarios. In George E. Marcus, editor, *Corporate Futures: The Diffusion of the Culturally Sensitive Corporate Form*, pages 141–176. University of Chicago Press, 1998.

[41] Anind Dey. *Providing Architectural Support for Building Context-Aware Applications.* PhD thesis, College of Computing, Georgia Tech, December 2000.

[42] Anind K. Dey and Gregory D. Abowd. Towards a better understanding of context and context-awareness. In *Proceedings of the CHI 2000 Workshop on The What, Who, Where, When, Why and How of Context-Awareness*, April 2000.

[43] Anind K. Dey, Gregory D. Abowd, Mike Pinkerton, and Andrew Wood. Cyberdesk: A framework for providing self-integrating ubiquitous software services. In *Proceedings of the ACM Symposium on User Interface Software and Technology*, pages 75–76, 1997.

[44] Steve Ditlea. The electronic paper chase. *Scientific American*, 285(5):50–55, November 2001.

[45] Thomas Dübendorfer. An extensible infrastructure and a representation scheme for distributed smart proxies of real world objects. Technical Report 359, ETH Zurich, April 2001.

[46] Ken Ducatel, Marc Bogdanowicz, Fabiana Scapolo, Jos Leijten, and Jean-Claude Burgelman. Scenarios for ambient intelligence in 2010. Technical report, The European Union IST (Information Society Technologies) Advisory Group, 2001.

[47] Jonas Eliasson and Mattias Lundberg. Road pricing in urban areas. Technical report, Vägverket, Swedish National Road Administration, January 2003.

[48] Stephen C. Farber, Robert Costanza, and Matthew A. Wilson. Economic and ecological concepts for valuing ecosystem services. *Ecological Economics*, 41(3):375–392, 2002.

[49] Glover T. Ferguson. Have your objects call my objects. *Harvard Business Review*, 80(6):138–144, 2002.

[50] Lilia Filipova and Peter Welzel. Reducing asymmetric information in insurance markets: Cars with black boxes. In *Proceedings of the 32nd Conference of the European Association for Research in Industrial Economics (EARIE)*, Porto, Portugal, September 2005.

[51] Elgar Fleisch. Von der Vernetzung von Unternehmen zur Vernutzung von Dingen. In Marcus Schögel, Torsten Tomczak, and Christian Belz, editors, *Roadm@p to E-Business – Wie Unternehmen das Internet erfolgreich nutzen*, pages 124–135. Thexis, 2002.

[52] Elgar Fleisch, Friedemann Mattern, and Stephan Billinger. Betriebswirtschaftliche Applikationen des Ubiquitous Computing –

Beispiele, Bausteine und Nutzenpotentiale. In Heinz Sauerburger, editor, *Ubiquitous Computing*, number 229 in HMD – Praxis der Wirtschaftsinformatik, pages 5–15. dpunkt.verlag, February 2003.

[53] Christian Floerkemeier. EPC-Technologie – vom Auto ID Center zu EPCglobal. In Elgar Fleisch and Friedemann Mattern, editors, *Das Internet der Dinge – Ubiquitous Computing und RFID in der Praxis.* Springer-Verlag, 2005.

[54] Christian Floerkemeier, Matthias Lampe, and Thomas Schoch. The smart box concept for ubiquitous computing environments. In *Proceedings of sOc'2003 (Smart Objects Conference)*, pages 118–121, Grenoble, France, May 2003.

[55] Christian Floerkemeier and Friedemann Mattern. Smart playing cards – enhancing the gaming experience with rfid. In *Proceedings of the Third International Workshop on Pervasive Gaming Applications - PerGames 2006 at PERVASIVE 2006*, Dublin, Ireland, May 2006.

[56] Christian Floerkemeier and Frank Siegemund. Improving the effectiveness of medical treatment with pervasive computing technologies. UbiHealth 2003: The 2nd International Workshop on Ubiquitous Computing for Pervasive Healthcare Applications, October 2003.

[57] Errol C. Friedberg. *The Writing Life of James D. Watson.* Cold Spring Harbor Laboratory Press, 2005.

[58] Masaaki Fukumoto and Mitsuru Shinagawa. Carpetlan: A novel indoor wireless(-like) networking and positioning system. In Michael Beigl, Stephen Intille, Jun Rekimoto, and Hideyuki Tokuda, editors, *Proceedings of the 7th Conference on Ubiquitous Computing (UbiComp 2005)*, volume 3660 of *LNCS*, pages 1–18, Tokyo, Japan, September 2005. Springer-Verlag.

[59] David Gay, Phil Levis, Rob von Behren, Matt Welsh, Eric Brewer, and David Culler. The nesc language: A holistic approach to networked embedded systems. *Proceedings of Programming Language Design and Implementation (PLDI) 2003*, June 2003.

[60] Hans-Werner Gellersen, Michael Beigl, and Holger Krull. The mediacup: Awareness technology embedded in an everyday ob-

ject. In Hans-Werner Gellersen, editor, *Proceedings of the 1st International Symposium on Handheld and Ubiquitous Computing (HUC'99)*, volume 1707 of *LNCS*, pages 308–310, Karlsruhe, Germany, 1999. TecO, University of Karlsruhe, Springer-Verlag.

[61] Hans-Werner Gellersen, Albrecht Schmidt, and Michael Beigl. Adding some smartness to devices and everyday things. In *Proceedings of the IEEE Workshop on Mobile Computing Systems and Applications (WMCSA)*, pages 3–10, Los Alamitos, CA, USA, December 2000.

[62] Neil Gershenfeld. *When Things Start to Think*. Henry Holt, 1999.

[63] Sharon Ghamari-Tabrizi. *The Worlds of Herman Kahn: The Intuitive Science of Thermonuclear War*. Harvard University Press, 2005.

[64] Theodore J. Gordon and Olaf Helmer-Hirschberg. Report on a long-range forecasting study. Technical Report P-2982, RAND Corporation, Santa Monica, CA, USA, 1964.

[65] Jonathan Grudin. Desituating action: Digital representation of context. *Human-Computer Interaction*, 16(2-4):269–286, 2001.

[66] Peter Hadfield. Smart cars steer round traffic jams. *New Scientist*, 154(2079):22, April 1997.

[67] Jessica Heesen, Christoph Hubig, Oliver Siemoneit, and Klaus Wiegerling. Leben in einer vernetzten und informatisierten Welt – Context-Awareness im Schnittfeld von Mobile und Ubiquitous Computing. Technical Report 05, German Center of Excellence 627: NeXuS – Spatial World Models for Mobile Context-Aware Applications, March 2005.

[68] Sumi Helal, Steven Moore, and Balaji Ramachandran. Drishti: An integrated navigation system for visually impaired and disabled. In *Proceedings of the 5th IEEE Symposium on Wearable Computers (ISWC'01)*, Zurich, Switzerland, 2001.

[69] Kahn Herman and Anthony J. Wiener. *The Year 2000: A Framework for Speculation on the Next Thirty-Three Years*. Macmillan, 1967.

[70] Ralf Guido Herrtwich. Fahrzeuge am Netz. In Friedemann Mattern, editor, *Total vernetzt – Szenarien einer informatisierten Welt*, pages 63–83. Springer-Verlag, 2003.

[71] Jeffrey Hightower and Gaetano Borriello. Location systems for ubiquitous computing. *IEEE Computer*, 34(8):57–66, August 2001.

[72] Jeffrey Hightower and Gaetano Borriello. Location systems for ubiquitous computing. *IEEE Computer*, 34(8):57–66, August 2001.

[73] Jason Hill, Robert Szewczyk, Alec Woo, Seth Hollar, David Culler, and Kristofer Pister. System architecture directions for networked sensors. *ACM SIGPLAN Notices (Special Issue: Proceedings of Architectural Support for Programming Languages and Operating Systems)*, 35(11):93–104, November 2000.

[74] Lorenz Hilty, Siegfried Behrendt, Mathias Binswanger, Arend Bruinink, Lorenz Erdmann, Jürg Fröhlich, Andreas Köhler, Niels Kuster, Claudia Som, and Felix Würtenberger. Das Vorsorgeprinzip in der Informationsgesellschaft – Auswirkungen des Pervasive Computing auf Gesundheit und Umwelt. TA 46/2003, TA-Swiss, 2003.

[75] Lorenz Hilty, Siegfried Behrendt, Mathias Binswanger, Arend Bruinink, Lorenz Erdmann, Jürg Fröhlich, Andreas Köhler, Niels Kuster, Claudia Som, and Felix Würtenberger. The precautionary principle in the information society – effects of pervasive computing on health and environment. TA 46e/2005, TA-Swiss, 2005.

[76] Lorenz M. Hilty, Andreas Köhler, Fabian Von Schéele, Rainer Zah, and Thomas Ruddy. Rebound effects of progress in information technology. *International Journal of Technology Assessment and Ethics of Science*, 4(1):19–38, March 2006.

[77] Lorenz M. Hilty, Claudia Som, and Andreas Köhler. Assessing the human, social, and environmental risks of pervasive computing. *Human and Ecological Risk Assessment*, 10(5):853–874, 2004.

[78] Matthias Horx. *Das Zukunfts-Manifest. Aufbruch aus der Jammerkultur.* Econ Tb., 2000.

[79] Michael Hoy. Categorizing risks in the insurance industry. *The Quarterly Journal of Economics*, 97(2):321–336, 1982.

[80] Andreas Hub, Joachim Diepstraten, and Thomas Ertl. Design and development of an indoor navigation and object identification system for the blind. *ACM SIGACCESS Accessibility and Computing*, (77–78):147–152, September 2003.

[81] Christoph Hubig. *Wirkliche Virtualität - Medialitätsveränderungen der Technik und der Verlust der Spuren*, pages 1–20. Bielefeld: Transcript, January 2005.

[82] Food Marketing Institute. Supermarket facts: Industry overview 2006. http://www.fmi.org/facts_figs/superfact.htm, 2007.

[83] Yogesh V. Joshi. Information visibility and its effect on supply chain dynamics. Master's thesis, MIT, 2000.

[84] Herman Kahn. *Thinking About the Unthinkable.* Horizon Press, 1950.

[85] Herman Kahn. *On Thermonuclear War.* Princeton University Press, 1960.

[86] Joseph M. Kahn, Randy H. Katz, and Kristofer S. J. Pister. Next century challenges: Mobile networking for "smart dust". In *Proceedings of the ACM/IEEE International Conference on Mobile Computing and Networking (MobiCom 99)*, pages 271–278, Seattle, WA, USA, August 1999.

[87] Tim Kindberg and John Barton. A web-based nomadic computing system. *Computer Networks*, 35(4):443–456, March 2001.

[88] Tim Kindberg, John Barton, Jeff Morgan, Gene Becker, Debbie Caswell, Philippe Debaty, Gita Gopal, Marcos Frid, Venky Krishnan, Howard Morris, John Schettino, Bill Serra, and Mirjana Spasojevic. People, places, things: Web presence for the real world. *Mobile Networks and Applications*, 7:365–376, 2002.

[89] Art Kleiner. *The Age of Heretics: Heroes, Outlaws, and the Forerunners of Corporate Change.* Doubleday, 1996.

[90] Art Kleiner. The man who saw the future. *strategy+business*, Spring 2003.

[91] Raymond J. Kopp and V. Kerry Smith. Valuing natural assets: The economics of natural resource damage assessment. RFF Press, 1993.

[92] David M. Kreps. *A Course in Microeconomic Theory.* Princeton University Press, 1990.

[93] John Krumm, Gerry Cermak, and Eric Horvitz. Rightspot: A novel sense of location for a smart personal object. In Anind K. Dey, Albrecht Schmidt, and Joseph F. McCarthy, editors, *Proceedings of the 5th Conference on Ubiquitous Computing (UbiComp 2003)*, volume 2864 of *LNCS*, pages 36–43, Seattle, WA, USA, October 2003. Springer-Verlag.

[94] Marc Langheinrich. The 5th dimension – building blocks for smart infrastructures. Position Paper, September 2000.

[95] Marc Langheinrich. Privacy by design – principles of privacy-aware ubiquitous systems. In Gregory D. Abowd, Barry Brumitt, and Steven A. Shafer, editors, *Proceedings of the Third International Conference on Ubiquitous Computing (UbiComp 2001)*, number 2201 in LNCS, pages 273–291, Atlanta, USA, 2001. Springer-Verlag.

[96] Marc Langheinrich. A privacy awareness system for ubiquitous computing environments. In Gaetano Borriello and Lars Erik Holmquist, editors, *Proceedings of the Fourth International Conference on Ubiquitous Computing (UbiComp 2002)*, number 2498 in LNCS, pages 237–245, Gothenburg, Sweden, September 2002. Springer-Verlag.

[97] Marc Langheinrich. *Personal Privacy in Ubiquitous Computing – Tools and System Support.* PhD thesis, ETH Zurich, Zurich, Switzerland, May 2005.

[98] Marc Langheinrich, Vlad Coroamă, Jürgen Bohn, and Michael Rohs. As we may live – real-world implications of ubiquitous computing. Technical Report, April 2002.

[99] Marc Langheinrich and Friedemann Mattern. Digitalisierung des Alltags. Was ist Pervasive Computing? *Aus Politik und Zeitgeschichte*, (42):6–12, October 2003.

[100] Hau L. Lee, V. Padmanabhan, and Seungjin Whang. The bullwhip effect in supply chains. *MIT Sloan Management Review*, 38(3):93–102, 1997.

[101] Lawrence Lessig. *Code and Other Laws of Cyberspace.* Basic Books, New York, USA, 1999.

[102] Todd Litman. Distance-based vehicle insurance as a TDM strategy. *Transportation Quarterly*, 51(3):119–138, 2007.

[103] Jack M. Loomis, Reginald G. Golledge, and Roberta L. Klatzky. Navigation system for the blind: Auditory display modes and guidance. *Presence: Teleoperators and Virtual Environments*, 7(2):193–203, 1998.

[104] Konrad Lorincz, David J. Malan, Thaddeus R.F. Fulford-Jones, Alan Nawoj, Antony Clavel, Victor Shnayder, Geoffrey Mainland, Steve Moulton, and Matt Welsh. Sensor networks for emergency response: Challenges and opportunities. *IEEE Pervasive Computing*, 3(4):16–23, 2004.

[105] Robert Lucky. Everything will be connected to everything else. *Connections. IEEE Spectrum*, March 1999.

[106] Christopher Lueg. On the gap between vision and feasibility. In Friedemann Mattern and Mahmoud Naghshineh, editors, *Proceedings of the 1st International Conference on Pervasive Computing (Pervasive 2002)*, volume 2414 of *Lecture Notes in Computer Science*, pages 45–57, Zurich, Switzerland, August 2002. Springer-Verlag.

[107] David Lyon. Facing the future: Seeking ethics for everyday surveillance. *Ethics and Information Technology*, 3(3):171–180, July 2001.

[108] Thomas Maeder. What barbie wants, barbie gets. *Wired Magazine*, 10(1):6, 2002.

[109] Eleonora Barbieri Masini and Javier Medina Vasquez. Scenarios as seen from a human and social perspective. *Technological Forecasting and Social Change*, 65(1):49–66, September 2000.

[110] Friedemann Mattern. The vision and technical foundations of ubiquitous computing. *Upgrade*, 2(5):2–6, October 2001.

[111] Friedemann Mattern. Vom Verschwinden des Computers – Die Vision des Ubiquitous Computing. In Friedemann Mattern, editor, *Total vernetzt – Szenarien einer informatisierten Welt*, pages 1–41. Springer-Verlag, 2003.

[112] Friedemann Mattern. Die technische Basis für das Internet der Dinge. In Elgar Fleisch and Friedemann Mattern, editors, *Das Internet der Dinge – Ubiquitous Computing und RFID in der Praxis*, pages 39–66. Springer-Verlag, 2005.

[113] Friedemann Mattern. Hundert Jahre Zukunft: Visionen zum Computer- und Informationszeitalter. In Friedemann Mattern, editor, *Die Informatisierung des Alltags – Leben in smarten Umgebungen*, pages 351–419. Springer, Berlin Heidelberg New York, 2007.

[114] Friedemann Mattern. Allgegenwärtige Informationsverarbeitung – Technologietrends und Auswirkungen des Ubiquitous Computing. In Alexander Rossnagel, Tom Sommerlatte, and Udo Winand, editors, *Digitale Visionen - Zur Gestaltung allgegenwärtiger Informationstechnologien*, pages 3–29. Springer, Berlin Heidelberg New York, April 2008.

[115] Florian Michahelles, Timo Ahonen, and Bernt Schiele. A-life – increasing survival chances in avalanches by wearable sensors. *3rd International Workshop on Smart Appliances and Wearable Computing*, 2003.

[116] Alex Mihailidis, Brent Carmichael, Jennifer Boger, and Geoff Fernie. An intelligent environment to support aging-in-place, safety, and independence of older adults with dementia. UbiHealth 2003: The 2nd International Workshop on Ubiquitous Computing for Pervasive Healthcare Applications, October 2003.

[117] Gordon Moore. Cramming more components onto integrated circuits. *Electronics*, (38):114–117, April 1965.

[118] Deirdre Mulligan and Ari Schwartz. Your place or mine?: Privacy concerns and solutions for server and client-side storage of personal information. In *Proceedings of the Tenth Conference on Computers, Freedom and Privacy*, pages 81–84, 2000.

[119] Matthias Oberholzer. *Strategische Implikationen des Ubiquitous Computing für das Nichtleben-Geschäft im Privatkundensegment der Assekuranz.* PhD thesis, Basel University, Switzerland, 2003.

[120] United Kingdom Office of National Statistics. Volume of retail sales at 2000 prices, seasonally adjusted. http://www.statistics.gov.uk/downloads/theme_economy/SDM28_May_2007.pdf, 2007.

[121] Torsten Olderog and Bernd Skiera. The benefits of bundling strategies. *Schmalenbach Business Review*, 1:137–160, 2000.

[122] Jason Pascoe. Adding generic contextual capabilities to wearable computers. In *Proceedings of the 2nd International Symposium on Wearable Computers ISWC*, pages 92–99, 1998.

[123] Kimberly Patch. Radio tags give guidance. www.trnmag.com/Stories/2003/092403/Radio_tags_give_guidance_092403.html, 2003.

[124] Arthur Cecil Pigou. *The Economics of Welfare.* Macmillian, London, 1st edition, 1920.

[125] Christoph Plüss. Smarter Fahrtenschreiber – Erfassung und Bearbeitung von Fahrzeugdaten zur verursachergerechten Verteilung der Verkehrskosten. Master's thesis, ETH Zurich, Institute for Pervasive Computing, 2004.

[126] Joseph Polastre, Robert Szewczyk, Cory Sharp, and David Culler. The mote revolution: Low power wireless sensor network devices. Stanford, CA, USA, aug 2004.

[127] Lucienne Rey. publifocus – road pricing. Technical report, TA-Swiss, July 2004.

[128] Gill Ringland. *Scenario Planning. Managing for the Future.* Wiley & Sons, 1997.

[129] Ron L. Rivest. Whither information security? SCS Distinguished Lecture, March 2001.

[130] Yvonne Rogers. Moving on from Weiser's Vision of Calm Computing: Engaging UbiComp Experiences. In Paul Dourish and

Adrian Friday, editors, *Proceedings of the 8th Conference on Ubiquitous Computing (UbiComp 2006)*, volume 4206 of *LNCS*, pages 404–421, Orange County, CA, USA, 2006. Springer.

[131] Kay Römer. *Time Synchronization and Localization in Sensor Networks.* PhD thesis, ETH Zurich, Zurich, Switzerland, June 2005.

[132] Kay Römer, Thomas Schoch, Friedemann Mattern, and Thomas Dübendorfer. Smart identification frameworks for ubiquitous computing applications. *Wireless Networks*, 10(6):689–700, December 2004.

[133] Joseph Rose. Closet buddy: Dressing the visually impaired. In *Proceedings of the 44th Annual Southeast Regional Conference*, pages 611–615, Melbourne, FL, USA, March 2006.

[134] Alexander Roßnagel. *Rechtswissenschaftliche Technikfolgenforschung – Umrisse einer Forschungsdisziplin.* Nomos Verlagsgesellschaft, Baden-Baden, 1993.

[135] Felix Röthenbacher. World explorer for visually impaired people. Master's thesis, ETH Zurich, Institute for Pervasive Computing, October 2003.

[136] Thomas Rudolph and Markus Schweizer. Problemfelder und Herausforderungen im Handel. Technical Report, University of St. Gallen, December 2001.

[137] Mahadev Satyanarayanan. Interview: M. Satyanarayanan on Mobile and Pervasive Computing. *IEEE Distributed Systems Online*, 2(6), 2001.

[138] Bill Schilit, Norman Adams, and Roy Want. Context-aware computing applications. In *Proceedings of the 1st IEEE Workshop on Mobile Computing Systems and Applications (WMCSA'94)*, pages 85–90, Santa Cruz, CA, US, 1994.

[139] Erwin Schrödinger. *What is Life? The Physical Aspect of the Living Cell.* Cambridge University Press, 1944.

[140] Peter Schwartz. *The Art of the Long View: Planning for the Future in an Uncertain World.* Bantam Doubleday Dell Publishing Group, 2 edition, 1996.

[141] Ludwig Siegele. How about now? A survey of the real-time economy. *The Economist*, 362(8257):3–18, January 2002.

[142] Ryan Singel. Nun terrorized by terror watch. *Wired News*, September 2005.

[143] Bernd Skiera and Martin Spann. Preisdifferenzierung im Internet. In Marcus Schögel, Torsten Tomczak, and Christian Belz, editors, *Roadm@p to E-Business – Wie Unternehmen das Internet erfolgreich nutzen*, pages 270–284. Thexis, 2002.

[144] Robert Sloss. Das drahtlose Jahrhundert. In Arthur Brehmer, editor, *Die Welt in 100 Jahren*, pages 27–48. Verlagsanstalt Buntdruck, Berlin, 1910.

[145] Yehuda Sonnenblick. An indoor navigation system for blind individuals. In *Proceedings of the 1998 CSUN Techology and Disabled Persons Conference*, Los Angeles, CA, USA, 1998.

[146] "Der Standard" staff. Wien droht doppelt so viel Autoverkehr. *Der Standard – Austrian daily newspaper*, April 24th-25th, 2004.

[147] The Economist staff. Driven to radicalism. *The Economist*, June 11th:33–34, 2005.

[148] The Economist staff. Jam yesterday. *The Economist*, June 11th:14, 2005.

[149] Richard M. Stallman. Can you trust your computer? In *Free Software, Free Society: Selected Essays of Richard M. Stallman*, pages 115–118. GNU Press, 2002.

[150] Joanna Stavins. Price discrimination in the airline market: The effect of market concentration. *The Review of Economics and Statistics*, 83(1):200–202, February 2001.

[151] Angelika Steinmüller and Karlheinz Steinmüller. *Visionen 1900 2000 2100. Eine Chronik der Zukunft.* Rogner & Bernhard, 1999.

[152] Lars A. Stole. Price discrimination and imperfect competition. In Mark Armstrong and Robert Porter, editors, *Handbook of Industrial Organization*, volume 3, chapter 8. North-Holland, Amsterdam, 2007.

[153] Robert Szewczyk, Alan Mainwaring, Joseph Polastre, John Anderson, and David Culler. An analysis of a large scale habitat monitoring application. In *Proceedings of the 2nd International Conference on Embedded Networked Sensor Systems*, pages 214–226, Baltimore, MD, USA, 2004.

[154] Clemens Szyperski. *Component Software – Beyond Object-Oriented Programming.* Addison-Wesley, 1999.

[155] Kees van der Heijden. *Scenarios: The Art of Strategic Conversation.* Wiley, 2 edition, 2005.

[156] Hal Varian. Price discrimination and social welfare. *American Economic Review*, 75(4):870–875, September 1985.

[157] Hal Varian. Price discrimination. In *Handbook of Industrial Organization*, volume 1, chapter 10, pages 597–654. North-Holland, 1989.

[158] Frank Verhoven. International price discrimination in the european car market. *RAND Journal of Economics*, 27(2):240–268, 1996.

[159] Jules Verne. *De la Terre à la Lune.* Pierre-Jules Hetzel, 1865.

[160] Pierre Wack. Scenarios: Shooting the rapids. *Harvard Business Review*, 63(6):139–150, November 1985.

[161] Pierre Wack. Scenarios: Uncharted waters ahead. *Harvard Business Review*, 63(5):73–89, September 1985.

[162] Roy Want, Andy Hopper, Veronica Falcao, and Jonathan Gibbons. The active badge location system. *ACM Transactions on Information Systems*, 10(1):91–102, January 1992.

[163] Mark Weiser. The computer for the 21st century. *Scientific American*, 265(3):66–75, September 1991. Reprinted in IEEE Pervasive Computing, 1(1), Jan.-Mar. 2002, pp. 19–25.

[164] Mark Weiser. Some computer science issues in ubiquitous computing. *Communications of the ACM*, July 1993.

[165] Mark Weiser and John Seely Brown. The coming age of calm technology. In Peter J. Denning and Robert M. Metcalfe, editors, *Beyond Calculation: The Next Fifty Years of Computing.* Springer-Verlag, 1998.

[166] Organic light-emitting diode (August 20th, 2007, 08:32 UTC). In *Wikipedia: The Free Encyclopedia*. Wikimedia Foundation, en.wikipedia.org/wiki/Organic_light-emitting_diode, 2007.

[167] Price Discrimination (June 30, 2007, 15:59 UTC). In *Wikipedia: The Free Encyclopedia*. Wikimedia Foundation, en.wikipedia.org/wiki/Price_discrimination, 2007.

[168] Ultra-wideband (April 26th, 2008, 11:23 UTC). In *Wikipedia: The Free Encyclopedia*. Wikimedia Foundation, en.wikipedia.org/wiki/Ultra-wideband, 2008.

[169] Sara Willliams and Charlie Kindel. The component object model: A technical overview. Technical report, Microsoft Corporation, 1994.

Curriculum Vitae

Vlad-Constantin Coroamă

Personal Data

Date of Birth	February 2, 1974
Birthplace	Bucharest, Romania
Citizenship	Romanian

Education

1980–1988	*Hermann-Oberth-Schule*, German-speaking elementary school, Bucharest, Romania.
1988–1992	*Hermann-Oberth-Gymnasium*, German-speaking high-school, Bucharest, Romania.
June 1992	high-school degree (Abitur).
1992–1994	Study of Electrical Engineering at the *Polytechnic University Bucharest*, Romania.
1994–2000	Study of Computer Science (Informatik) at the *Technical University of Darmstadt*, Germany.
March 2000	Master of Science in Computer Science (Diplom-Informatiker).
2000–2008	Ph.D. Student at the Department of Computer Science, *ETH Zürich*, Switzerland.

Employment

1995–1998	Student software developer at the *Fraunhofer Institute for Computer Graphics*, Darmstadt, Germany.
1998–2000	Software architect and developer at the Cryptography and Computer Algebra Group, *TU Darmstadt*, Germany.
2000–2006	Research Assistant at the Department of Computer Science, *ETH Zurich*, Switzerland.
2006–2007	Technology consultant at *Helbling AG*, Aarau, Switzerland.
2007–2008	Senior researcher at the Computer Networks Laboratory of the *University of São Paulo*, Brazil.