

HANDBOOK OF RESEARCH ON

# Ubiquitous Computing Technology for Real Time Enterprises



MAX MÜHLHÄUSER & IRYNA GUREVYCH

TP274-62  
H236.2

# Handbook of Research on Ubiquitous Computing Technology for Real Time Enterprises

Max Mühlhäuser  
*Technische Universität Darmstadt, Germany*

Iryna Gurevych  
*Technische Universität Darmstadt, Germany*



E2009003731

**Information Science**  
**REFERENCE**

**INFORMATION SCIENCE REFERENCE**

Hershey • New York

Acquisitions Editor: Kristin Klinger  
Development Editor: Kristin Roth  
Editorial Assistants: Ross Miller, Deborah Yahnke  
Senior Managing Editor: Jennifer Neidig  
Managing Editor: Sara Reed  
Copy Editors: Alana Bubnis, Erin Meyer  
Typesetter: Michael Brehm  
Cover Design: Lisa Tosheff  
Printed at: Yurchak Printing Inc.

Published in the United States of America by  
Information Science Reference (an imprint of IGI Global)  
701 E. Chocolate Avenue, Suite 200  
Hershey PA 17033  
Tel: 717-533-8845  
Fax: 717-533-8661  
E-mail: [cust@igi-global.com](mailto:cust@igi-global.com)  
Web site: <http://www.igi-global.com>

and in the United Kingdom by  
Information Science Reference (an imprint of IGI Global)  
3 Henrietta Street  
Covent Garden  
London WC2E 8LU  
Tel: 44 20 7240 0856  
Fax: 44 20 7379 0609  
Web site: <http://www.eurospanonline.com>

Copyright © 2008 by IGI Global. All rights reserved. No part of this publication may be reproduced, stored or distributed in any form or by any means, electronic or mechanical, including photocopying, without written permission from the publisher.

Product or company names used in this set are for identification purposes only. Inclusion of the names of the products or companies does not indicate a claim of ownership by IGI Global of the trademark or registered trademark.

#### Library of Congress Cataloging-in-Publication Data

Handbook of research on ubiquitous computing technology for real time enterprises / Max Muhlhauser and Iryna Gurevych, editors.  
p. cm.

Summary: "This book combines the fundamental methods, algorithms, and concepts of pervasive computing with current innovations and solutions to emerging challenges. It systemically covers such topics as network and application scalability, wireless network connectivity, adaptability and "context-aware" computing, information technology security and liability, and human-computer interaction"--Provided by publisher.

Includes bibliographical references and index.

ISBN-13: 978-1-59904-832-1 (hardcover)

ISBN-13: 978-1-59904-835-2 (ebook)

1. Ubiquitous computing--Handbooks, manuals, etc. 2. Real-time data processing--Handbooks, manuals, etc. I. Muhlhauser, Max. II. Gurevych, Iryna.

QA76.5915.H35 2007

004--dc22

2007032050

#### British Cataloguing in Publication Data

A Cataloguing in Publication record for this book is available from the British Library.

All work contributed to this book set is original material. The views expressed in this book are those of the authors, but not necessarily of the publisher.

*If a library purchased a print copy of this publication, please go to <http://www.igi-global.com/reference/assets/IGR-eAccess-agreement.pdf> for information on activating the library's complimentary electronic access to this publication.*

# Editorial Advisory Board

Erwin Aitenbichler  
*Technische Universität Darmstadt, Germany*

Hidir Aras  
*University of Bremen, Germany*

Gerhard Austaller  
*Technische Universität Darmstadt, Germany*

Alexander Behring  
*Technische Universität Darmstadt, Germany*

Dirk Bradler  
*Technische Universität Darmstadt, Germany*

Martin Faust  
*University of Bremen, Germany*

Iryna Gurevych  
*Technische Universität Darmstadt, Germany*

Andreas Hartl  
*Technische Universität Darmstadt, Germany*

Melanie Hartmann  
*Technische Universität Darmstadt, Germany*

Andreas Heinemann  
*Technische Universität Darmstadt, Germany*

Marc Herrlich  
*University of Bremen, Germany*

Jussi Kangasharju  
*Technische Universität Darmstadt, Germany*

Tobias Klug  
*Technische Universität Darmstadt, Germany*

Tobias Limberger  
*Technische Universität Darmstadt, Germany*

Fernando Lyardet  
*Technische Universität Darmstadt, Germany*

Max Mühlhäuser  
*Technische Universität Darmstadt, Germany*

Christof Müller  
*Technische Universität Darmstadt, Germany*

Daniel Oberle  
*SAP Research, Germany*

Andreas Petter  
*Technische Universität Darmstadt, Germany*

Robert Porzel  
*University of Bremen, Germany*

Sebastian Ries  
*Technische Universität Darmstadt, Germany*

Dirk Schnelle  
*Technische Universität Darmstadt, Germany*

Daniel Schreiber  
*Technische Universität Darmstadt, Germany*

Julian Schröder-Bernhardi  
*Technische Universität Darmstadt, Germany*

Jürgen Steimle  
*Technische Universität Darmstadt, Germany*

Joachim Steinmetz  
*Technische Universität Darmstadt, Germany*

Stefan-Georg Weber  
*Technische Universität Darmstadt, Germany*

Markus Weimar  
*Technische Universität Darmstadt, Germany*

Torsten Zesch  
*Technische Universität Darmstadt, Germany*

## Foreword

### WHAT'S IN A NAME IF IT IS ALL IN THE GAME?

When reading through the manuscript of this novel volume I was struck by the heroic attempt of the editors to position their book as a holistic approach to the subject of ubiquitous computing. I found their strong stand especially striking in this respect with respect to the use of nomenclature in the domain of ubiquitous computing. The editors acknowledge that there are many different notions presented in the literature addressing similar concepts as that of ubiquitous computing, but they argue that all these notions should be considered as a single approach to the topic of the disappearing computer. More specifically, the editors refuse to identify and describe the borderlines between different notions such as ubiquitous computing, pervasive computing, and ambient intelligence, following their strong conviction that it makes not much sense to quarrel about thin borderlines between major overlapping fields as their exploration is still open to a large extent.

As a convert to the concept of ambient intelligence for almost ten years now I must admit that I continuously have felt the need in the past to explain these differences in an attempt to mark the borderlines. Evidently, most of these notions, which were developed during the late nineties of the past century, are rooted in the early ideas expressed by the late Mark Weiser, who was dreaming of a world that would be flooded with embedded devices, note pads, and electronic dust, which would soon become feasible as a result of the remarkable advances in the manufacturing of semiconductor devices and micro-systems. However, the developments that have been achieved over the past ten years have shown that there can be no doubt about the question whether or not Mark's dream will come true; it surely will. The remaining question however is related to the issue of which form it will take and how it can be configured in such a way that society and its participants maximally benefit from it. On the other hand, some of the innovation directions have changed in the meantime, which has opened new venues for research. Great inventions, such as ambient atmospheres through distributed solid-state lighting devices, virtual environments applying 3D interactive words such as Second Life, and ultimately "The Internet of Things" have made the discussion about the differences between the various notions artificial and esoteric. More interesting therefore is the question how far the advances in this domain have stretched the boundaries of what is currently feasible. And again the editors deserve a compliment as they have addressed this question in a most original way. Their S.C.A.L.E. classification provides a simple and most practical reference model for the description of the relevant topics in the field of ubiquitous computing. Furthermore, they have succeeded in combining in the present book a most remarkable collection of research results representative of the advances in this domain. The many high-quality contributions reflect the scholarship and expertise of their authors. The book is definitely a mandatory reading for anyone who is professionally active in the field of ubiquitous computing, as it can be seen as a landmark approach to the description of the advances in this domain.

After more than ten years of breakthrough developments, ubiquitous computing can now live up to its expectation that it can change peoples' lives for the better through the promise of the disappearing computer.

Finally, I would like to thank the editors for providing me with the insight that the true progress achieved in our field if investigation is not reflected by names we attribute to our inventions, but merely by the games we play with it.

*Emile Aarts*

**Emile Aarts** holds an MSc and a PhD degree in physics. For more than 20 years he has been active as a research scientist in computing science. Since 1991 he has held a teaching position at the Eindhoven University of Technology as a part-time professor of computing science. He also serves on numerous scientific and governmental advisory boards. He holds a part-time position of senior consultant with the Center for Quantitative Methods in Eindhoven, The Netherlands. Aarts is the author of 10 books and more than 150 scientific papers on a diversity of subjects including nuclear physics, VLSI design, combinatorial optimization and neural networks. In 1998 he launched the concept of ambient intelligence, and in 2001 he founded Philips' HomeLab. His current research interests include intelligent systems and interaction technology.

## Foreword

As computing has become more and more an integral part of our daily business and personal lives, the trend of ubiquitous computing will transform the way in which businesses will work and collaborate. The well-known fact of a huge community of users on the Internet (1,100 million users as of March 2007) will be complemented by at least one order of magnitude higher (10,000 millions of artificial users) instantiated by machines, sensors and things connected to the Internet. More precise data will be generated and accumulated that enable completely new business scenarios for the future. The fact of being connected to that universe of human users and artificial users will speed up decisions in business (real-time enterprise) and enable those who can master the infrastructure and the application services on top of the infrastructure to be more competitive than others. Application fields from logistics to e-health, from supply-chain management and manufacturing to public security will benefit from the fact that the “Internet of Things” and the “Internet of People” converge using an “Internet of Services” architecture.

I would like to congratulate Professor Mühlhäuser and Dr. Gurevych for their comprehensive overview of ubiquitous computing scenarios, real world examples and architectural blueprints that combine the various elements into insights into the vision of how the virtual world will interact with the physical world. I would also like to thank my colleagues from SAP Research and the SAP senior executives who have been supporting the research program of “smart items,” which has produced many excellent results over the last eight years that are also reflected in this book.

*Joachim Schaper*

**Joachim Schaper** received his Diploma (1988) and PhD (1995) from the Technical University of Karlsruhe. Since 1989, he worked for Digital Equipment Corp. in their European Research Center, CEC Karlsruhe. He became the manager of that center, which in turn became part of SAP AG Corporate Research (1999). In 2001, Schaper took over additional responsibilities as a founding manager of the Corporate Research Groups at SAP Labs France and SAP Africa. From 2003 to 2005, he managed the SAP Research Center in Palo Alto and a research group in Montreal. A vice president of EMEA, Schaper is responsible for all research activities of SAP in Europe, Middle East, and Africa, reporting to the head of corporate research and to the executive board. His research interests comply with the topics investigated in the SAP research groups on e-learning, smart items, mobile computing, and technology for application integration and advanced customer interfaces.

# Preface

## ABSTRACT

*The preface provides an introduction to and a definition of ubiquitous computing as a computer science field and relates it to the concept of real time enterprises. We describe the main challenges in ubiquitous computing and introduce the S.C.A.L.E. classification employed to organize the contents of the book. Finally, recommendations about using the book as a reference and teaching resource are given.*

## OUTLINE AND SUBJECT OF THIS BOOK

On the next couple of pages, we first want to provide an introduction to and a definition of ubiquitous computing (UC)—both as a scientific domain and as a technology area—and relate it to the concept of real time enterprises. We first define the scope of UC as a domain—in particular as covered in the present book. The question is raised whether UC is a research field in its own right; we also explain the required trade-off between breadth and depth of coverage concerning the contents of this book. The S.C.A.L.E. classification is introduced as probably the first attempt to provide a canonical structure and organization of the area. The present preface thus gives the reader a better idea about why this book is structured in a particular way and why it covers the range of topics selected by the editors. A “reader’s digest” is provided, both as an overview of the chapters provided in this book and as a guide for readers with different backgrounds and interests.

So far, no single book exists which sufficiently covers ubiquitous computing in a holistic way. Many UC books restrict themselves to combinations of middleware, networking, and security. However, UC has lately extended beyond this focus and even beyond purely technical issues. In particular, understanding current developments in the field requires knowledge about pertinent algorithms and concepts in artificial intelligence and human-computer interaction. All-in-one reference books covering the foundations of ubiquitous computing and the areas mentioned above are missing; therefore, researchers, practitioners, and academics typically use collections of papers from the respective conferences and individual chapters from the books emphasizing a single area. This approach does not provide the target audience with a coherent view of the field. Also, the presentation of materials often presumes too much knowledge about the related topics.

As we will substantiate later, real time enterprises (RTE) are a key application area for UC. In fact, the authors of this book will show that RTE is *the* first application domain when it comes to the economic advantages of UC. Therefore, RTE can be considered a key driving force for large-scale deployment of UC technology. The last part of this book describes a selection of pilot projects and trend analyses concerning UC applied to RTE, provided by *SAP Research*. These chapters will strongly support the

above-mentioned arguments. Note that most UC concepts and technologies described in this book are not restricted to RTE. Rather, readers should consider RTE as the first widespread, commercially successful area - they should also consider our RTE related examples as an aid for a better understanding of the value and usage of UC.

## DEFINING UC

According to the Oxford English Dictionary, the word *ubiquitous* has two meanings: the first meaning is *(seemingly) present, everywhere simultaneously*, and the second meaning is *often encountered*. The seminal work by Weiser (1991) introduced the term ubiquitous computing, stating that it:

“represents a powerful shift in computation, where people live, work, and play in a seamlessly interweaving computing environment. Ubiquitous computing postulates a world where people are surrounded by computing devices and a computing infrastructure that supports us in everything we do.”

Well, then what is *ubiquitous computing*?

One approach to an answer is the rough division of computer science into three consecutive eras:

1. Era number one was that of *mainframe computers*, where one computer was used by many users (1:N)
2. Era number two is about to end: the era of *personal computers* (PC), where one computer was used (owned) by one user (1:1), and
3. The third, dawning era is one in which many computers surround a single user (N:1)—almost anytime and anywhere.

Based on this approach, ubiquitous computing is nothing else than the third era cited above, that is, it is equivalent to the “Post-PC era.”

We all experience this dawning era and realize that it brings about a proliferation of computers, with desktop PCs, laptops, PDAs and cell phones just being examples. The anytime-anywhere availability of computers indicates a shift away from pure desktop—and thereby, “isolated, full-attention”—computing to mobile (or rather, nomadic)—and thereby “integrated, shared-attention” computing. The term *integrated* alludes to the fact that the new computers interact with—or are even perceived as “part of”—their environment; the attribute *shared-attention* emphasizes the fact that users do not devote themselves to “sessions at the computer,” but rather conduct an activity during which they “happen to interact with a computer, too.”

Accordingly, the computers that surround us fall into two categories:

- Some are *worn or carried* in the sense that they move along with us, and
- Some are *encountered* in the sense that they are either part of the environment of our respective whereabouts, or worn or carried again, namely by people whom we meet.

The first category covers devices denoted as wearables or portable devices, respectively. The second category of devices encountered in the environment shifts from traditional PCs—which were somewhat alien to the environment and general-purpose in nature—to computers that are perceived as an integral part of the environment and that are rather special-purpose in nature (see the next section for concrete examples). More precisely speaking, the boundaries between general-purpose computers (formerly: servers and PCs) and special-purpose computers (formerly: microcontrollers) become blurred. On the

one hand, computers become ever cheaper and can be dedicated to specific tasks. Power constraints and other resource limitations for portable and unattended computers increase this trend towards dedicated devices and favor “right-sizing.” On the other hand, the required flexibility, adaptability, sophistication, and maintainability suggest devices that can be easily re-programmed. As a negative side effect, the latter trend introduces the curses of general-purpose computers—vulnerability (e.g., in the IT security sense) and limited reliability due to restricted maturity of devices with fast innovation cycles (cf., the fact that cell phones now have a far greater tendency to “crash,” since they are based on full operating systems and software applications).

One is tempted to define UC simply as the era of portable and embedded specialized computers. However, embedded computers have existed for decades and have *already* become ubiquitous: they have become indispensable in washing machines and VCRs, and up to a hundred or even more of them are on duty in modern automobiles. *New* is the fact that these embedded devices are now:

- Networked, that is, cooperation enabled, Internet-enabled, and
- More flexible in terms of both maintenance/evolution and adaptation.

(Rather) new is also the insight that neither PCs—with their still unnatural and non-intuitive interfaces, interaction devices, and modes-of-operation (updates, re-boots, ...)—nor embedded devices—with their increasing “featurism” (increasing sophistication that average users don’t learn how to exploit)—scale well up to a world where hundreds of them would surround a user. Quantum leaps are required in terms of ease-of-use if ubiquitous computers are not to become a curse.

In summary, we can turn the rough definition “UC is the Post-PC era” into the following more elaborate one:

*Ubiquitous computing is the dawning era of computing, in which individuals are surrounded by many networked, spontaneously yet tightly cooperating computers, some of them worn or carried, some of them encountered on the move, many of them serving dedicated purposes as part of physical objects, all of them used in an intuitive, hardly noticeable way with limited attention.*

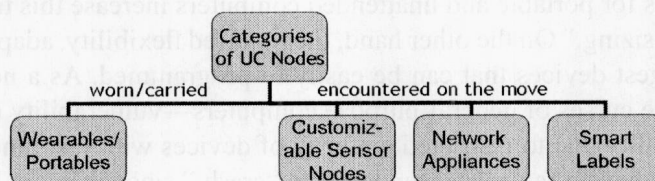
As a link to upcoming sections, readers should note that the definition given above buries two crucial issues:

1. Truly *integrative* cooperation: the path from mere connectivity of the “networked UC nodes” to true cooperation is long and arduous. Rapid deployment of all kind of networked sensors, appliances, labels, and so forth does not by itself lead to a meaningful “whole.”
2. As to the last line of the definition above, it was already mentioned that a quantum leap in usability and dependability is required. We want to call this requirement a need for *humane* computing henceforth.

## THE SIGNIFICANCE OF UC

The reader may want to get a feeling about the spatial and temporal significance of UC. By “spatial” we mean the spectrum from niche technologies to technologies that deeply influence an era. In this respect, it should soon become clear that UC is going to deeply mark our society over the years to come. By “temporal” we mean how visionary, that is, far off, UC actually is. In this respect, the quick answer, which we are going to substantiate further in the following, is as follows:

Figure 1. Overview of UC nodes



1. On one hand, *UC is already a reality* in the sense that the computer-based devices carried and encountered by users are already—and increasingly—networked.
2. On the other hand, *UC is still a big challenge* with respect to the required quantum leaps in *integrative* cooperation, as mentioned further previously, (“the whole must become way more than its parts”) and *humane* behavior.

A few examples are cited for the first aspect, by providing a very selective list of four important categories of “UC nodes” in the global UC network:

1. **Wearables and portable devices** (“networked computers worn or carried” as mentioned in the UC definition), such as handhelds for warehouse picking, washable computer jackets, or companions like the so called *lovegetties* introduced in Japan almost ten years ago. These pocket size devices store their user’s profile, that is, dating interests, and beep or vibrate when a “compatible” person—carrying a lovegetty—appears. In contrast to portable device, the term *wearable* denotes a degree of integration with a piece of clothing or clothing accessory that goes beyond that of mobile computing devices, up to a degree where the “computer nature” is hardly noticed by the user. MIT is known for influential research in the field ([www.media.mit.edu/wearables](http://www.media.mit.edu/wearables)). The examples given illustrate that this category of UC nodes is far larger than a simple extrapolation from the common representatives, that is, PDAs, cell phones, MP3 players, laptops, and so forth. Another example, namely body sensors, illustrates the possible seamless integration of “UC nodes worn or carried” and “UC nodes encountered.” Body sensors are useful, for example, for activity monitoring (a basis for better context-aware applications) and health condition monitoring.

The last three of these categories are refinements of the “networked computers encountered on the move” from the definition of UC given in the last section (with exceptions like body sensors).

2. **Customizable sensor nodes** like the so-called *motes* developed by UC Berkeley and Intel. As opposed to traditional sensors, these nodes come with a fully programmable microprocessor and micro operating system (called *TinyOS* for *motes*), a variety of sensor options, and low energy networking. Mass production of these customizable nodes is intended to drive down cost such that easy customization and easy assembly of nodes into *sensor networks* are set to simplify application development. The Intel-funded company Crossbow ([www.xbow.com](http://www.xbow.com)) is commercializing motes. Companies like the German start-up *Particle Computer* ([www.particle-computer.de](http://www.particle-computer.de)) support the IEEE sensor network standard 802.15.4 known as *ZigBee* ([www.zigbee.org](http://www.zigbee.org)). In 2003, UC Berkeley built the first single-chip successor to motes nicknamed *Spec*. They represent a major step towards almost invisible sensor networks with large quantities of nodes, often coined as *smart dust*.
3. **Networked appliances**, also called *Internet appliances* or *smart appliances*, which are mainly perceived by their users as tools, machines, devices, furniture, and so forth, rather than computers.

Apart from ever cheaper and more powerful embedded computers, affordable and energy-aware wireless technology is a key enabler not only for hand-held appliances, but also for fixed installations—which are much easier to deploy if network cabling can be spared on and if embedded “clients” can be easily combined with back office servers or gateways inaccessible to the public. One example is *smart vending machines*: the company USA technology ([www.usatech.com](http://www.usatech.com)) developed solutions for supporting online transactions (credit card payment, etc.) and for transmitting various information (fill status or out-of-stock events, out-of-change events, defects, customer behaviour patterns, etc.) to the operating agency. One step further, vending of physical and of digital goods start to converge: Coke vending machines are combined with vending of cell phone credits, ring tones, MP3 music, and so forth.

4. **Smart labels**, which identify physical objects and creatures vis-à-vis an IT-based system. Radio frequency identifiers (*RFIDs*) represent one important technology; *active badges* (e.g., for employees) denote a possible use case. Such labels may be thought of as Web addresses that serve as a “link” to data describing details about the identified object or person. Therefore, their on-board storage and processing capabilities can be very limited, even non-existent. In the application domain of *product identification*, RFIDs are set to complement (and later, replace) barcodes. Relevant barcode standards include the Universal Product Code UPC and (as a superset) the European Article Number EAN. For RFIDs, both are intended to be replaced by EPCglobal’s Electronic Product Code EPC ([www.epcglobalinc.org](http://www.epcglobalinc.org)), which may contain a serial number in addition to the article number.

These few categories of ubiquitous computing nodes illustrate that specialized networked computers—most of them “hidden” in physical objects—are indeed penetrating the world.

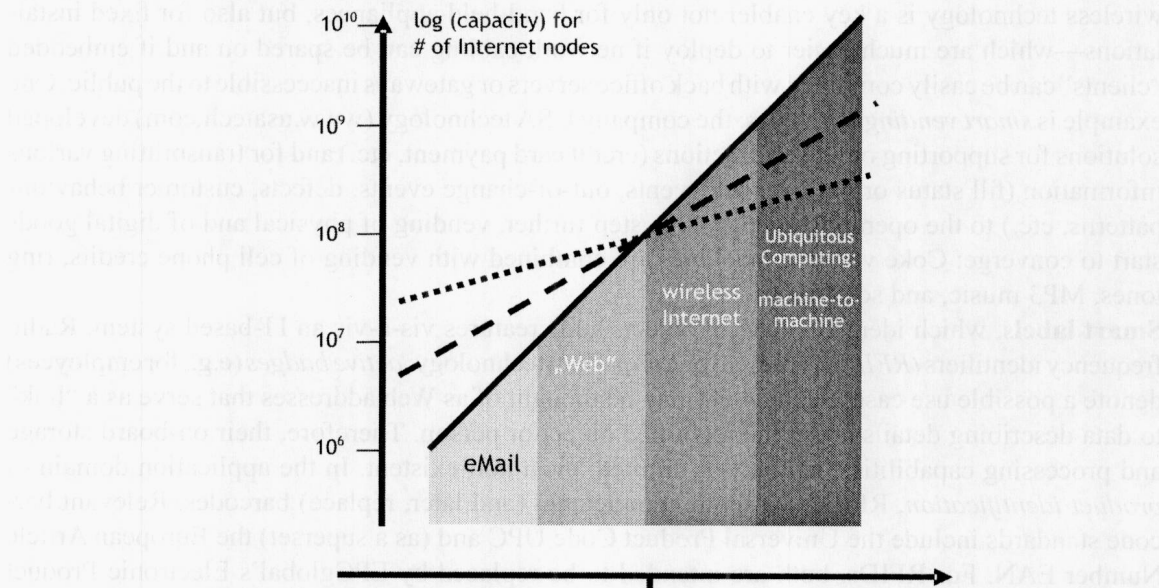
## Machine-to-Machine Communication

The example categories cited above illustrate yet another phenomenon of key importance: the advent of large volume machine-to-machine communication. Throughout the history of the Internet, its exponential growth has been boosted by communication with or among people. This development contradicted what protagonists of distributed parallel processing had envisioned, namely a boosting importance of distributed systems as replacements for parallel and super computers. While this use case for distributed systems—and thereby for the Internet—continues to play a non-negligible role (cf., recent advancements in Grid computing), it has never been, and probably will never be a reason for “exploding numbers of Internet nodes.” Three “waves of Internet usages” have driven and are driving this explosion to date:

1. “E-mail” as a means for people-to-people communication, the major reason for the Internet to grow up to the order of 10 million nodes.
2. “The Web” as a major use case of people-to-machine communication. Under its increasing importance, the Internet passed the 100 million node mark
3. “The wireless Internet,” which—after flops like WAP-based cell phones—currently drives the convergence of the Internet with cell phones (cf., UMTS in Korea), MP3 players, and other portable devices. This wave contributes heavily to the hit of 1 billion Internet nodes around 2008.

As we approach the order of magnitude of the world population, one might expect the Internet (number of nodes) to enter a phase of sub-exponential growth. Yet the predicted shift from people-to-machine to machine-to-machine communication is likely to lead to a continuation of the exponential growth: sensor

Figure 2. Growth rates: Internet vs. CPU power



networks, networked appliances, and smart labels communicate autonomously and forward real-world events through the net.

A second important effect is illustrated in Figure 2: the growth rate of the Internet has constantly exceeded that of CPU power—measured, for example, in terms of the time that it takes for key indicators to double, like the number of Internet nodes or the processor speed (18 to 20 months for the latter one according to “Moore’s Law”).

This difference in growth rate has consequences for the Internet: the *relative* cost of distributed versus local processing decreases. Note, however, that this is only true for the wired Internet, for which the aggregated throughput per second (at least nationwide for the U.S.) and the typical bandwidth are all growing at roughly the same pace as the number of Internet nodes. Two reasons suggest more conservative predictions for wireless nodes: on one hand, there is no long term indication yet that the typical bandwidth will keep pace with the other Internet growth indicators mentioned above. Despite recent boosts in WLAN bandwidth, physical limits and frequency scarcity remain tough obstacles. But even if bandwidth keeps pace, wireless nodes are set to become more and more independent from the power line—and for mass deployment use cases like sensor nodes, even from battery replacement. A question mark must be placed here at least until revolutionary battery technology hits a breakthrough.

In summary, the reader should now understand that UC is the *key technology* that will deeply influence our society for three reasons:

1. UC describes *the* next era of computing. Since we live in the information (i.e., computer) society, the influence will be at least as pervasive as that of computer today. As a side remark, virtually every domain of computer science or IT is potentially impacted. This makes it difficult to be selective and concise for the present book—but not impossible, as we will show.
2. UC has potential impact on every facet of our lives. Computing is no longer “what we do when we sit at the computer” nor “what is encapsulated/hidden deep inside VCRs, and so forth.”
3. UC is inevitable and “impossible” at the same time: the components are already developed and massively deployed, consult the four categories of UC nodes described. Since UC use cases are

becoming increasingly profitable, for example, the replacement of barcodes with RFIDs, the industry will push the use of UC technology. Nevertheless, the two top-level challenges remain “integrative cooperation” and “humane computing.” Both must be solved in order for UC to become the envisioned helpful anytime-anywhere technology and not a nightmare.

## THE CHALLENGES OF UC

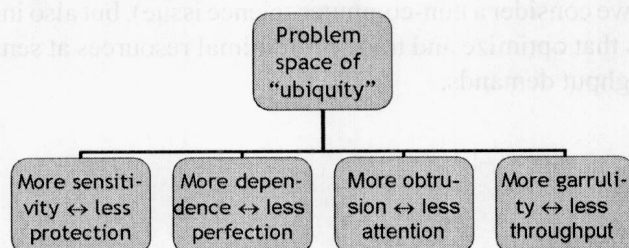
As described in the last sentences above, UC is inevitable and probably very influential, even able to change our society—and it buries unresolved problems. Therefore, UC can be considered one of the biggest challenges of our times—an insight that led to the present book, for instance.

### Conflicting Developments Due to Ubiquity of Networked Computers

The above sentences have also pointed out two top-level challenges in UC: “integrative cooperation” and “humane computing.” Before we try to describe and define these challenges in more detail, as a basis for the structure and content of this book, we want to mention one more viewpoint of what has been already discussed: we will try to describe the upcoming “ubiquity” of networked computers as a problem space that is becoming increasingly challenging due to conflicting developments on four levels (see Figure 3):

1. **More sensitivity ↔ less protection:** As UC penetrates more and more areas of daily life and work life, more and more sensitive data, but also processes and activities, of privacy critical and liability critical nature depend on computers. As a conflicting development, the established IT security solutions are not fully viable in machine-to-machine scenarios, in particular if a subset of these machines acts autonomously as substitutes for humans or organizations. Further aggravation of the situation is due to the ever widening gap between the “in principle” availability of appropriate security measures and their acceptance, that is, application by users, even more so since IT related concepts—for example, of trust—do not easily match the concepts of trust that users are acquainted with from their real-life experience. Obviously, IT security research—in a broad sense, including trust models and other issues—must be stressed and focused on UC.
2. **More dependence ↔ less perfection:** if many actions and aspects of daily life are supported and controlled by ubiquitous computers, then these actions and aspects (and in turn the average human) become dependent on these computers. Obviously, this dependence becomes as ubiquitous as the computers, and is not restricted to desktop work anymore. As an increasingly critical source of

Figure 3. Overview of conflicting developments



conflict, ubiquitous computing relies on cheap components and close to zero maintenance, a context in which the failure of nodes becomes a regular event as opposed to an exception. Furthermore, overall system dependability declines rapidly with the number of components involved under the assumption that the per-component dependability (and therefore, failure rate) remain constant and that all considered components contribute a crucial portion to the system. Obviously, research on system dependability must emphasize “over-provisioning” approaches where, for instance, a considerable number of components can fail without harmful consequences for system dependability. Complex biological and social systems exhibit this property and can serve as examples.

3. **More obtrusion ↔ less attention:** users are expected to be surrounded by ubiquitous networked computers for many purposes. In principle, each additional computer means an additional user interface. As a source of conflict, the move from desktop use to everyday use means that users have to share attention to computers with attention to their daily lives (human communication, driving, manual work, etc.). Altogether, this issue suggests a dramatic focus on “humane computing” as has been stated several times already. We will see below that such research must be understood in a broader sense than classical human-computer interaction (HCI) research.
4. **More garrulity ↔ less throughput:** the market for sensors is predicted to explode and the replacement of barcodes by RFIDs will lead to zillions of wireless nodes (in the longer term, after intermediate phases where only containers and expensive goods will be tagged). These two effects alone show how dramatically the number of wireless data paths around us will grow—and feed the wired network in turn. One may object to calling this development dramatic, since machine-to-machine communication can often be restricted to short binary messages, whereas computer-to-human communication tends to involve verbose data, media streams in particular. However, increasing performance of computer vision and voice recognition will increase the bandwidth famine on the sensor-to-backend path. Again, one may object that computer vision may be carried out locally and that only reduced information will be sent (intermediate results, i.e., recognized “features” of the scene observed, or the final result, i.e., a “semantic” description of what was recognized). However, computer vision is compute intensive and therefore not likely to be executed entirely on the sensor side. Increased sensor-to-backend bandwidth demands can therefore be expected at least for this kind of sensor.

The critical factor in this fourth dimension is the fact that sensor nodes must be built under extreme energy constraints, such as a grain size battery determining the lifetime of the entire sensor. Since wireless transmission tends to be the most energy hungry part of a sensor and since this hunger increases considerably as bandwidth is increased, we can state that the throughput of many UC nodes will remain far below that of an average PC even today. Similar arguments apply to smart labels: the widely used RFID tags are passive, that is, receive their transmitting energy from the reader (receiver). Portable readers are crucial for many application scenarios, and the required energy is a function of both bandwidth and distance—yet increased reading distance is supposed to be a key advantage over barcodes. All in all, there is a large demand for advancements not only in the energy/bandwidth tradeoff (and in battery capacity, of course, which we consider a non-computer science issue), but also in distributed architectures and distributed algorithms that optimize and trade off minimal resources at sensor nodes (and other UC nodes) and minimal throughput demands.

## Integrative Cooperation and Humane Computing

The preceding observations provide further insight into the challenges of UC, which we described as “integrative cooperation” and “humane computing” on the top level. We now want to go one level deeper and divide each of these two challenges into two or three sub-issues.

As to “integrative cooperation,” we want to distinguish two issues:

1. **Scalability:** We discussed several times that UC furthers the explosion of the Internet and that even many local networks (such as smart dust and RFID-tagged products in a warehouse) rather represent big assemblies of nodes (in particular in relation to the network bandwidth to be expected). This means that highly scalable solutions are required. The history of the Internet (especially in comparison to less successful networks) is a history of extreme scalability, and all standard textbooks about the Internet state that scalability was and is a key success factor. This applies even more to ubiquitous computing as the future nature of IT systems, depending on the Internet. Note that scalability is to be understood in a broad sense, as will be described further below.
2. **Connectivity:** Again, this term is a brief notion for a broad issue. We have identified scalability as a broad issue in its own right, but everything else that relates to the need to integrate UC nodes and to make them cooperate will be summarized under *connectivity*. In particular, tiny (“dumb,” or at least special-purpose) nodes must be integrated into a meaningful whole based on spontaneous discovery and interaction. New, that is unforeseen, components or components that vary over time or usage context must be able to enter into meaningful cooperation with existing nodes. We will see later in the book that such connectivity issues are addressed and solved somewhat differently for different kinds of UC components today; for example, differently for software services than for smart labels. Various chapters of the book will reflect these differences that exist today. When it comes to the vision of “intelligent behavior” of the integrated, cooperating whole, we must deal with approaches that are also related to the scalability issue and to the ease-of-use issues treated below. This indicates that our classification (as always) is good for understanding the topic space, but not “ideal” in the sense that any issue could be classified as belonging to exactly one category (an issue that any taxonomy has to live with).

The residual three issues represent sub-challenges of “humane computing”:

3. **Adaptability:** A key to dramatically reduced cognitive load (required in order to counter the “more obtrusiveness, less attention” dilemma described) is found in highly adaptive systems that interact “optimally”—with respect to the use case, user, situation, device or resource constraints, and so forth. A great deal of research in this respect has been carried out lately under the term “context-aware computing;” many of the respective activities tried to adapt software and particularly user interfaces to data retrieved from sensors with the users’ locations being a number one context category. Other context types, like activities carried out by the user or data “hidden” in their files, are also important, but their integration into UC solutions is less well understood. Even less emphasized and understood today is adaptation to the users’ state of knowledge.
4. **Liability:** As discussed with the issue “more sensitivity, less protection,” IT security issues must be revisited under UC requirements and constraints. Moreover, the use of UC technology in everyday life makes UC-based physical and digital components an integral part of our society—and consequently of our economy. A majority of UC components or services will not be available for free. Even if they are free of charge to the end-user, someone will have to pay for their development

and execution. This means that UC services will have to respond to a number of “market rules,” for instance: (i) users will want to make sure that “they get what they paid for,” (ii) providers will want to make sure that “they get paid for what they provide,” and (iii) users will want to be able to truly compare offerings based on prices and a variety of parameters describing the associated “value” of a UC service. Providers will want to be able to offer their “values” to users appropriately, and so forth. Since all these issues are tightly linked to financial, and thereby to legal issues, they are closely intertwined with the security issues mentioned above; altogether, we will summarize the respective issues as “liability.”

5. **Ease-of-use:** Under this term, we want to emphasize HCI aspects in the more narrow sense. Maybe the single most important change brought about by UC in this respect is the proliferation of interaction devices and, thereby, modalities. In order to understand this issue, one has to note that throughout the history of computer science, software engineers and programmers were exposed to a single major interaction paradigm at any given time. It started with “holes in paper,” that is, punch tapes and punch cards. Later teletype like devices came, followed by full screen terminals like the IBM 3270 and the DEC VT100, both being de facto industry standards. Finally, more than 15 years ago, Windows-based UIs and the WIMPS metaphor (windows, icons, menus, pointers, scrollbars) were introduced. In the UC world of special-purpose nodes, there is no longer any single dominating interaction paradigm: cooperative-use devices like interactive walls, speech-and-graphics devices, soft-key-based cell phones, physical-interaction appliances, and many more must be coped with. Accordingly, multimodal user interfaces—supposed to cater for this variety—are becoming crucial. At the same time, the “classical” graphical (or rather, hands-and-eyes) interfaces are currently experiencing rapid innovation and “voice” is recognized as a most natural way of interaction, in particular for users whose hands and eyes are busy with “real world activities.” These and other issues must be addressed under “ease-of-use” in a UC related book.

The five sub-challenges treated above—scalability, connectivity, adaptability, liability, and ease-of-use—can be abbreviated by concatenating the first letter of each sub-challenge, leading to S.C.A.L.E. We will describe below how they represent the backbone structure for the present book—and a suggested structure for the emerging discipline.

## THE FOCUS ON REAL TIME ENTERPRISES

This section addresses two questions: “What is a real time enterprise?” And, “Why and how does the present book emphasize this topic?”

First of all, in preparation of a “definition,” one must understand that in many domains software has become more and more sophisticated over the years. Whatever domain is supported—it is represented as a machine-readable “digital model” as part of the software. For instance, *virtual reality* applications for the automotive industry, such as “digital mockups,” contain not only three-dimensional representations of a yet-to-be-built car but also the materials’ characteristics, laws of physics used to model the dynamics—to a degree that supports “virtual crash tests”—and much more. The same applies to enterprise application integration (EAI) software (at the core of real time enterprise approaches): almost the entire enterprise is “digitally modeled,” including human resources, workflows, a whole spectrum of financial issues, product lifecycle management, customer relationships, and much more.

Despite the increasing sophistication of “digital enterprise models,” the gap between this digital model and the “real enterprise” remains hard to bridge: events and changes in the real enterprise (goods arriv-