Appears in: Journal of Human and Ecological Risk Assessment, Vol. 10, No. 5, pp. 763-786, October 2004.

Living in a World of Smart Everyday Objects – Social, Economic, and Ethical Implications

Jürgen Bohn, Vlad Coroamă, Marc Langheinrich, Friedemann Mattern[*], Michael Rohs
Department of Computer Science, Institute for Pervasive Computing, ETH Zurich (Swiss Federal Institute of Technology), Haldeneggsteig 4, 8092 Zurich, Switzerland

ABSTRACT

Visions of pervasive computing and ambient intelligence involve integrating tiny microelectronic processors and sensors into everyday objects in order to make them "smart." Smart things can explore their environment, communicate with other smart things, and interact with humans, therefore helping users to cope with their tasks in new, intuitive ways. Although many concepts have already been tested out as prototypes in field trials, the repercussions of such extensive integration of computer technology into our everyday lives are difficult to predict. This article is a first attempt to classify the social, economic, and ethical implications of this development.

Key Words: Pervasive computing, ubiquitous computing, technological risk, privacy, ethical implications, social implications, economy.

INTRODUCTION

The increasing miniaturization of computer technology will, in the foreseeable future, result in processors and tiny sensors being integrated into more and more everyday objects, leading to the disappearance of traditional PC input and output media such as keyboards, mice, and screens. Instead, we will communicate directly with our clothes, watches, pens, and furniture – and these objects will communicate with each other and with other people's objects (Mattern 2004).

More than 10 years ago, Mark Weiser foresaw this development and described it in his influential article "The Computer for the 21st Century" (Weiser 1991). Weiser coined the term "ubiquitous computing," referring to omnipresent computers that serve people in their everyday lives at home and at work, functioning invisibly and unobtrusively in the background and freeing people to a large extent from tedious routine tasks. In its 1999 vision statement, the European Union's Information Society Technologies Program Advisory Group (ISTAG) used the term "ambient intelligence" in a similar fashion 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" (Ahola 2001). Nowadays, the notion of "pervasive computing" is often used in an equivalent way.

The vision of a future filled with smart and interacting everyday objects offers a whole range of fascinating possibilities. For example, parents will no longer lose track of their children, even in the busiest of crowds, when location sensors and communications modules are sewn into their clothes. Similar devices attached to timetables and signposts could guide blind people in unknown environments by "talking" to them via a wireless headset (Coroama and Röthenbacher 2003). Tiny communicating computers could also play a valuable role in protecting the environment, for example as sensors the size of dust particles that detect the dispersion of oil spills or forest fires. Another interesting possibility is that of linking any sort of information to everyday objects, allowing for example future washing machines to query our dirty clothes for washing instructions.

[*] Corresponding author: Tel: +41 1 632 05 36, Fax: +41 1 632 16 59; mattern@inf.ethz.ch

While developments in information technology never had the explicit goal of changing society, but rather did so as a side effect, the above-mentioned visions expressly propose to transform society by fully computerizing it. It is therefore very likely that this will have long-term consequences for our everyday lives and ethical values that are much more far-reaching than the Internet with all its discussions about spam e-mails, cyber crime, and child pornography. With its orientation towards the public as well as the private, the personal as well as the commercial, it aspires to create technology that will accompany us throughout our entire lives, day in and day out. And if Mark Weiser's vision of "invisible computing" actually materializes, we won't even notice any of it.

It seems to be clear that with these technical developments – pushed through largely unnoticed by the general public and extending quite rapidly into our everyday lives – unanticipated (if not unacceptable) standards could soon be set for the rest of our lives. In the following, we examine the driving factors behind the visions of pervasive computing and ambient intelligence – from a technical as well as an economic perspective – and we try to illustrate the social and ethical implications of a "smart world" that connects everything to everything else, where anywhere can potentially be contacted from anywhere else, and where everybody could conceivably interact with anybody (and anything) else.

TECHNOLOGY TRENDS

The driving force behind continuing technological progress in the field of information technology is the long-term trend in microelectronics: Moore's Law (Moore 1965), drawn up in the late 1960s by Gordon Moore and roughly stating that the power of microprocessors doubles about every 18 months, has held true with astonishing accuracy and consistency. A similarly high increase in cost-efficiency can be observed for some other technological parameters such as storage capacity and communications bandwidth. To put it another way, prices for microelectronic functionality with an equivalent amount of computing power are falling radically over time. This trend, which is expected to continue for at least another 15 years, means that computer processors and storage components will become much more powerful, smaller, and cheaper in the future, so that there will be an almost unlimited supply of them.

Even more important are the results of microsystem technology and nanotechnology. These could lead, for example, to flexible displays or electronic paper. Another interesting development is radio sensors that can report their readings within a few meters distance without an explicit energy supply – such sensors obtain the necessary energy from the environment or directly from the measuring process itself.

Electronic labels (so-called "smart labels" or RFID tags) also operate without their own energy supply. Depending on their construction, these are less than a square millimeter in area and thinner than a piece of paper. In some ways, this is a further development of the well-known anti-theft technology involving security gates in department stores. However, this is not just about the binary information "paid / stolen"; within milliseconds, several hundred characters could be read and written "wirelessly" up to a distance of a few meters (Finkenzeller 2003).

What is interesting about such remote-inquiry electronic markers is that they enable objects to be clearly identified and recognized, and therefore linked in real time to an associated data record held on the Internet or in a remote database. This ultimately means that specific data can be associated with any kind of object. If everyday objects can be uniquely identified from a distance and furnished with information, this opens up application possibilities that go far beyond the original purpose of automated warehousing or supermarkets without cashiers.

Significant advances have also been made in the field of wireless communications. Especially interesting are recent short-range communications technologies that require very little energy, making it possible to produce designs that are much smaller and cheaper than today's mobile phones. Intensive research is also being carried out into improved options for indicating the position of mobile objects. As well as increased accuracy (currently around ten meters for the GPS system), the aim is also to make the devices much smaller.

If you summarize these technology trends and developments – tiny, cheap processors with integrated sensors and wireless communications capability, attaching information to everyday objects, the remote identification of objects, the precise localization of objects, flexible displays based on polymers, and electronic paper – it becomes clear that the technological basis for a strange new world has been created: everyday objects that are in some respects “smart,” and with which we can even communicate under certain circumstances.

There are various ways of implementing such communication with things. As one example, imagine everyday objects such as furniture, packaged food, medication, clothing, or toys being equipped with an electronic label containing a specific Internet address as digital information. If you can then read this Internet address with a portable device just by pointing it at the object, this device can, independently and with no further assistance from the object in question, access and display the associated information from the Internet via the mobile phone network. The user has the impression that the object itself has “transmitted” the information, although in fact it has been supplied to the display device via the Internet. The information could be, for example, operating instructions, or cooking instructions for a ready-to-serve meal, or the information leaflet for medication. The details of what is displayed may depend on the “context” – for example, whether the user is a good customer and paid a lot of money for the product, whether he is over 18 years of age, what language he speaks, or his current location, – but also maybe whether he has paid his taxes on time…

The foreseeable technological developments will therefore add an additional new quality to everyday objects – these might be able not only to communicate with people and other “smart” objects, but also to discover where they are, which other objects are in their vicinity, and what has happened to them in the past, for example. Objects and devices could thus behave in a context-sensitive manner and appear to be “smart,” without actually being “intelligent.”

While technological advances such as miniaturization, increasing computing power, and wireless connectivity open up the possibility of new applications, critics (Adamowsky 2000; Araya 1995; Lucky 1999) argue that it is not yet clear how these possibilities are actually going to be put into practice: we are “brilliant on means, but pretty hopeless when it comes to ends” (Thackara 2001). However, this innovation dilemma – we may know how we can create incredible things, but we don’t know what needs they are supposed to meet – is only superficial. The following section describes potential economic benefits that ubiquitous-computing technology offers when it comes to industrial processes – benefits that will be a prominent driver for the proliferation of pervasive computing, perhaps even more than the above-mentioned technological progress itself.

PERVASIVE COMPUTING AND THE ECONOMY

“The most profound revolutions are not the ones trumpeted by pundits, but those that sneak in when we are not looking” (Weiser 1993, p.72). What Mark Weiser formulated over ten years ago accurately describes the current atmosphere surrounding the field of pervasive computing. While personal gadgetry in the form of smartphones and Internet fridges continues to bedazzle the press, industry has quietly 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 (Siegele 2002b), 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 extraction. Ultimately, the whole lifecycle of products, beginning with the “birth” of their components and ending with their complete consumption (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.

Staying on Track: Smart Products

Inventory management 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" (Lee *et al.* 1997). 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 are attenuated (Joshi 2000). 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 (Fleisch 2002). 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" (Ferguson 2002, p. 143). 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.

Anything, Anywhere: Innovative Business 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 (Ferguson 2002, p. 144) describes a future, maybe not that far off, in which the boundary between the real world and the world of information has become blurred. 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. 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" (Skiera and Spann 2002, p. 275) – economists call this "perfect price discrimination." Such individual prices have to be introduced carefully, however: the online-bookseller Amazon discovered that such perfection might not be universally appreciated, after their trial of individual DVD prices had to be suspended almost immediately due to massive customer criticism (Siegele 2002a; USA Today 2000).

The ultimate in shopping may be achieved when all the decision-making is removed from people, and things do the shopping themselves (see the section on social challenges below for a discussion of pervasive computing and consumer control). The business consultancy Accenture has already coined a phrase for this – "silent commerce" (Accenture Technology Labs 2001). 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" (Maeder 2002, p. 6).

Apart from such rather questionable ways of increasing sales revenues, however, the "smartness" of objects might also facilitate trading at a more fundamental level, by reducing information asymmetries (Akerlof 1970): goods could not only "talk" about their price, ingredients, and availability, but also provide a detailed history of their production, use, and repair. A used car could give a detailed list of the parts that have been replaced or repaired over the course of its lifetime, thus reducing the amount of trust the buyer must have in the seller. Organic food could provide consumers with a comprehensive history of its cultivation, fertilization, and processing, potentially increasing their willingness to pay a premium for it. Obviously, in order for consumers to trust this information, it would either require verification from trusted third parties, or some sort of technical means to prevent tampering (much like odometers today). However, by increasing the overall market transparency, both for new and used goods, an environment supporting trustworthy pervasive-computing systems could potentially increase overall market activities by reducing the uncertainty inherent in many transactions today.

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 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 – Buena Vista's Disney home video unit, for example, recently started selling DVDs that render themselves unplayable 48 hours after unpacking (Reuters 2003). 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 a ubiquitous-computing 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. However, as with the market transparencies described above, mechanisms would need to be in place in order to assure both parties that the sensor readings had not been tampered with.

No risk, no premium

Information asymmetry is not only a hindrance for trading, but also affects the risk management (moral hazard) of insurance companies. Typically, insurers use broad categories such as age or location in order to judge the risk involved in providing individual coverage, e.g., for health or car insurance. An ambient-intelligence environment, however, would allow for a much more fine-grained calculation of the individual risks involved. A smart car, for example, could provide detailed information about the driving style and parking habits of its owner, thus providing the insurer with a much better assessment of the likelihood of an accident or theft. For "safe" drivers, this would result in reduced insurance premiums as they would be exempted from paying for the insurer's uncertainty regarding their driving style.

Insurance rates could even become entirely dynamic, reducing or increasing the premium in a "pay-per-use" fashion in real time, e.g., during driving. Taken to the extreme, the current insurance rate, similar to today's indicators showing current gas consumption, could be calculated on the spot and displayed in an "insurance meter", based on the time of day, the route traveled, weather and road conditions, and of course driving style. Similar adjustments could be made for home owners' insurance (new furniture automatically registers its net worth when placed inside the home) or health insurance (smoking a single cigarette would increase the rate, a walk in the park decrease it). Even in the case of an actual insurance claim, smart objects could significantly lower the insurer's cost by observing the circumstances of an accident or by noticing which items were actually missing after a break-in.

Obviously, such detailed activity records would constitute a serious privacy hazard for many consumers. While technical remedies might be possible (e.g., data would be encrypted before being sent to the insurer, who would need an authorized key from the customer – or law enforcement – to read it in case of an accident), both the increased complexity and the limited usefulness to the insurer of such encrypted information would most likely favor the privacy-threatening alternatives. Even if customers had the choice of opting out, drivers who did not wish to pass their details on to their insurer would most likely have to pay a considerably higher premium, as the insurer's risk would be spread among fewer and fewer non-participating customers. This would not only limit the desired freedom of choice, particularly for low-income drivers, but also implicitly accelerate the move away from fixed rates to cheaper dynamic (but privacy-invading) solutions.

An Economy on Autopilot

Despite the number of potential economic advantages described above, there are also substantial risks involved when relying on ubiquitous-computing technologies for large parts of an economy. The increasing automation of economically relevant aspects and the exclusion of humans as decision makers could certainly become a cause for concern. Under "normal" circumstances, automated control processes increase system stability – machines are certainly much better than humans if they have to devote their whole attention to a particularly boring task. But situations that have not been anticipated in the software can easily have disastrous consequences if they are not directly controlled by humans (as tragic accidents involving airplane autopilots have shown in the past). Other problems might arise from the intricate interplay of several automated processes, which might quickly escalate into an unanticipated feedback loop that gets out of control. For example, the stock market crash of 1987 was partly caused by newly implemented trading software (Siegele 2002b), which was designed in such a way that it would trigger the sale of shares whenever a certain pattern appeared in the daily fluctuations of their prices. Since a large number of traders were using the same software, the first appearance of this pattern caused a flood of sales, which in turn reinforced the software's selling pattern and thus triggered the crash.

Another risk stems from the increased efficiency of an economy supported by ubiquitous-computing technologies. While detailed tracking and inventory systems, along with smart, "self-conscious" products, allow a company to trim production and stock keeping as much as possible, this lack of "slack" in the production process also increases the risk that unforeseen interruptions can have grave consequences. In the case of supply chain management, for example, the attenuation of the bullwhip effect permits a reduction in storage capacity. However, if all companies along the delivery chain drastically reduce their stocks, one small unexpected interruption in supply by the weakest member would immediately halt all production along the whole chain.

The automation and acceleration of the economy increases not only the potential for possible savings, but also influences the associated risks in these types of complex and sensitive systems. Nevertheless, the increased flexibility and control that a world of pervasive computing offers will certainly be an important incentive for businesses to drive its deployment in the near future. Therefore, it is important that the ambient-intelligence landscapes we build should be both reliable and socially acceptable. We will come back to these issues in a later section. For now, we want to take a closer look at what probably constitutes the most visible implication of an economy driven by pervasive-computing technologies, namely its threat to personal privacy.

PRIVACY IN PERVASIVE COMPUTING

Intelligent fridges, pay-per-use scenarios, and dynamic insurance rates paint a future in which all of our moves, actions, and decisions are recorded by tireless electronic devices, from the kitchen and living room of our homes to our weekend trips in our cars. Not surprisingly, many critics see this as "an attempt at a violent technological penetration of everyday life" (Araya 1995, p. 237), as the "feverish dream of spooks and spies – to plant a 'bug' in every object" (Talbott 2000) or even as "a project that aims at totality and, of course, verges on the totalitarian" (Adamowsky 2000).

By virtue of its very definitions, the vision of ubiquitous computing has the potential to create an invisible and comprehensive surveillance network, covering an unprecedented share of our public and private life: “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” (Lucky 1999, p. 19). And with the economic possibilities described above, such a comprehensive coverage seems more likely to be put into place: shopping without participating in comprehensive profiling, buying instead of renting items in a pay-per-use scheme, as well as fixed insurance schemes that do not constantly transmit information to insurers – all of this might become an expensive luxury for well-off citizens, while the population at large must trade in their privacy for increased productivity and market transparency. This might very well be self-inflicted: given the immediate economic returns of consumer loyalty programs or low insurance rates, the rather vague threats of future privacy violations are easily enough ignored. The following sections try to add a differentiated view to this problem, especially with respect to ubiquitous-computing technology, by first examining why personal privacy is desirable, describing when we feel that it has been violated, and then assessing how the deployment of future pervasive-computing systems will affect all that.

The Many Facets of Personal Privacy

Even though critics continue to argue that “all this secrecy is making life harder, more expensive, dangerous and less serendipitous” (Cochrane 2000), privacy is still predominantly seen as a fundamental requirement of any modern democracy (Rotenberg 2001). It is only when people are free to decide what to do with their lives, according to their interests and beliefs, and without fear of repression from their fellow citizens, that the necessary plurality of ideas and attitudes can develop that will prevent society being subjugated under a charismatic leader. Harvard law professor Lawrence Lessig (Lessig 1999) takes this requirement a step further and distinguishes between a number of motives for the protection of privacy in today’s laws and standards:

- *Privacy as Empowerment.* Seeing privacy mainly as informational privacy, its aim is to give people the power to control the publication and distribution of information about themselves (Westin 1967). A recent legal discussion surrounding this motivation revolved around the question of whether personal information should be seen as private property (which would entail the right to sell all or part of it as the owner sees fit) or as intellectual property (which would entitle the owner to certain inalienable rights, preventing him for example from selling the rights to his own name to anybody else).
- *Privacy as Utility.* From the viewpoint of the person involved, privacy can be seen as a utility providing more or less effective protection against nuisances such as unsolicited phone calls or emails. This view probably best follows Brandeis’ definition of privacy as “The right to be left alone,” where the focus is on minimizing the amount of disturbance for the individual (Warren and Brandeis 1890).
- *Privacy as Dignity.* Dignity not only entails being free from unsubstantiated suspicion (for example being the target of a wire tap, where the intrusion is usually not directly perceived as a disturbance), but also focuses on the equilibrium of information available between two people: as in a situation where you are having a conversation with a fully dressed person when you yourself are naked, any relationship where there is a significant information imbalance will make it much more difficult for those with less information about the other to keep their composure.
- *Privacy as a Regulating Agent.* Privacy laws and moral norms to that extent can also be seen as a tool for keeping checks and balances on the powers of a decision-making elite. By limiting information gathering of a certain type, crimes or moral norms pertaining to that type of information cannot be effectively controlled.

Depending on what kind of motives one assumes for preserving privacy, ubiquitous-computing technology can become the driving factor for changing the scope and impact of privacy protection as it exists today, and creating substantially different social landscapes in the future. This is because ubiquitous-computing technology influences two important design parameters relating to privacy: the ability to monitor and the ability to search (Lessig 1999).

Pervasive Computing and Surveillance

The conscious observation of the actions and habits of our fellow men is as old as mankind itself. However, observations using automated systems differ from our nosy neighbors in one important aspect: while in the “good old days”, anything out of the ordinary would attract the attention of our fellow citizens, it is now the ordinary, the everyday routine, that can be (and often is) the sole focus of tireless computerized monitoring. With ambient–intelligence technologies, today's monitoring capabilities can obviously be extended far beyond credit-card records, call logs, and news postings. Not only will the spatial scope of such monitoring activities be significantly extended in ambient-intelligence landscapes, but their temporal coverage will also greatly increase: starting from pre-natal diagnostics data stored on babies' hospital smart cards, to activity patterns in kindergarten and schools, to workplace monitoring and senior citizen's health monitoring.

Such comprehensive monitoring (or surveillance) techniques create new opportunities for what MIT professor emeritus Gary T. Marx calls border crossings: "Central to our acceptance or sense of outrage with respect to surveillance ... are the implications for crossing personal borders" (Marx 2001, p. 158). He goes on to define four such border crossings that form the basis for perceived privacy violation:

- *Natural Borders.* Physical borders of observability, such as walls and doors, clothing, darkness, and also sealed letters and phone conversations. Even facial expressions can represent a natural border against the true feelings of a person.
- *Social Borders.* Expectations with regard to confidentiality in certain social groups, such as family members, doctors, and lawyers. This also includes the expectation that your colleagues do not read personal fax messages addressed to you, or material that you leave lying around the photocopier.
- *Spatial or Temporal Borders.* The expectation by people that parts of their lives can exist in isolation from other parts, both temporally and spatially. For example, a previous wild adolescent phase should not have a lasting influence on the current life of a father of four, nor should an evening with friends in a bar influence his coexistence with work colleagues.
- *Borders due to Ephemeral or Transitory Effects.* This describes what is best known as a "fleeting moment," a spontaneous utterance or action that we hope will soon be forgotten or old pictures and letters that we put out in our trash. Seeing audio or video recordings of such events subsequently, or observing someone sifting through our trash, would violate our expectations of being able to have information simply pass away unnoticed or forgotten.

Putting pervasive-computing systems into place will most certainly allow far greater possibilities for such border crossings in our daily routines. Consider the popular vision of a wearable "memory amplifier" (Mayo 2001; Rhodes 1997), allowing its wearer to constantly record the events of her daily life in a lifetime multimedia diary. While at first sight such technology promises assistance to those of us who frequently tend to forget small details, it also has substantial consequences for our privacy borders stemming from ephemeral and transitory effects: any statement I make during a private conversation could potentially be played back if my conversation partner gave others access to her multimedia diary. Even if this information were never disclosed to others, the very thought of dealing with people who have a perfect memory (and thus would never forget anything) would probably have a considerable effect on our interpersonal relationships.

The problem of spatial and temporal borders, on the other hand, is well known from the field of consumer profiles. Although such profiles are often the subject of public debate, the social and legal attitudes towards them have, until now, been relatively relaxed. Consumer acceptance is also much higher than the frequent negative news coverage might indicate, mostly because their negative consequences are often perceived as being rather minor (such as unsolicited spam) compared to their advantages (e.g., monetary incentives in the form of discounts or rewards). However, there are well-known risks associated with profiles, and their adoption as the basis for pervasive computing would only exacerbate such problems. Besides the obvious risk of accidental leaks of information (O'Harrow 2001), profiles also threaten universal equality, a concept central to many constitutions, basic laws, and human rights, where "all men are created equal" (USA 1998). Even though an extensively customized ambient-intelligence future where I only get the information that is relevant to my profile holds great promise, the fact that at the same time a large amount of information might be deliberately withheld from me because I am not considered a valued recipient of such information, would constitute a severe violation of privacy for many people.

Applying ubiquitous-computing technology in areas with primarily social borders – for example where a close social group interacts only among itself, such as families (Nagel *et al.* 2001; Westerlund *et al.* 2001) or co-workers – might appear to alleviate some of the above concerns. Most participants would already share close relationships and tend to know a great deal about each other, without needing a system to compile a profile of their communication partner. Such systems, however, also raise the ante as to what type of information they handle. While a communication whiteboard for families may facilitate social bonding between physically and temporally separated members, it would also increase the risk of unwanted social border crossings by accidentally allowing Mum to read a message you left for your sister, or a visiting friend to be recorded in the house activity log even though you told grandma you would spend the weekend alone.

Natural borders, then, might be the easiest to respect when designing pervasive-computing systems. Here, the concept of surveillance is well known and usually fairly straightforward to spot, after all: if others are able to watch your actions behind closed doors, they are most certainly intruding on your privacy. Proponents of wearable computing systems often cite the fact that information could both be gathered and stored locally (i.e., on the user's belt, or within her shirt) as a turnkey solution for privacy-conscious technologists (Rhodes *et al.* 1999). Border crossings, however, are not only about who does something, but also what is happening. Even though a context-aware wearable system might keep its data to itself, its array of sensors nevertheless probe deep into our personal life, and the things they might find there could easily startle (and trouble) us, once such systems start anticipating our future actions and reactions. The feeling of having someone (or something) constantly looking over our shoulder and second-guessing us would certainly constitute a natural border crossing for most of us. And the temptation of law enforcement subpoenaing such information not only to determine your physical data (were you at the crime scene?) but also to guess your intentions (by assessing the data feed from our body sensors) would certainly motivate legislation that would make the deletion of such information a crime, just as recent legislation against cybercrime (Council of Europe 2001) does for computer log files.

The Power of Searching and Combining Information Bits

All these examples serve to show that ambient-intelligence landscapes, even when created for the greater good and with the best of intentions, will run a high risk of involuntarily threatening the personal borders that separate our private and public lives, simply because their monitoring capabilities will facilitate more of the border crossings described above. However, whether such crossings ultimately occur, given the opportunities created, will depend very much on the type of searching capabilities such pervasive-computing systems might offer.

Search technology is traditionally a topic in the fields of information retrieval or databases, rather than that of pervasive computing. However, the chances are high that such technology will be a basic building block of future ambient-intelligence landscapes, as many of the envisioned applications in the field require precisely these capabilities. An automated diary collecting 24/7 audio and video data would not be much use unless it was combined with powerful search and retrieval technology that allowed us to comb large amounts of data for very specific information. And the ability to combine different information sources, especially large, innocuous ones such as walking patterns or eating habits, is the backbone of any "smart" system, which must make the best use of a large variety of different sensor input to take decisions that make it appear to understand what is happening around us.

Having thus both monitoring and search capabilities at the very core of their architecture, pervasive-computing systems will very likely provide their developers, owners, and regulators with a significant tool for driving the future development of privacy concepts within society. Depending on the actual systems deployed, some of the motivating aspects of privacy as discussed above might become more or less prominent, thus influencing corresponding legal and social norms. However, as important as privacy is, it is merely the tip of the iceberg that constitutes the social implications of pervasive computing. In the following section we want to explore a bit more of the edges that lurk just below the waterline, where they seem to be of no immediate threat but where they could potentially become much more dangerous as the deployment of pervasive computing picks up speed.

SOCIAL CHALLENGES AND IMPLICATIONS

Life without computers is unimaginable for most of us today – embedded processors monitor the condition of high-risk patients around the clock, they control central heating in buildings, air conditioning in tunnels, and they safely guide airplanes between continents. The potential economic benefits of pervasive computing are certainly key factors for the further proliferation of information technology, such as novel indoor and outdoor positioning systems, ubiquitous communication platforms, and unobtrusive monitoring installations. This technology will form and shape the foundations of future ambient-intelligence landscapes.

As more and more objects and environments are being equipped with ubiquitous-computing technology, the degree of our dependence on the correct, reliable functioning of the deployed devices and microcomputers including their software infrastructures is increasing accordingly. Today, in most cases, we are still able to decide for ourselves whether we want to use devices equipped with modern computer technology (e.g., by choosing manual control for our central heating, or by deciding not to carry a mobile phone if we dislike the constant accessibility its usage implies). But in a largely computerized future, it might not be possible to escape from this sort of technologically induced dependence, which leads to a number of fundamental social challenges for future pervasive-computing systems. Privacy is just one of these challenges, though probably the most prominent one. However, the more thoroughly "computerized" our environment becomes, the more basic attributes of the world we live in will subtly change, such as its reliability, accessibility, and transparency. In the following, we attempt to identify these concerns and try to address additional ethical and social implications of future ambient-intelligence landscapes.

Reliability

The vision of pervasive and ubiquitous computing describes systems that work completely in the background, discreetly and unobtrusively helping us to carry out our tasks. Since our needs and circumstances can change over time, such systems must be able to adapt themselves dynamically to the current situation. In doing so, one crucial basic requirement is reliability in the broadest sense of the word. In addition to ensuring dependability from a technological point of view, a complex and highly dynamic system must also remain manageable and controllable, and must retain the ability to predict (and, to a certain extent, verify) that the system is behaving correctly:

- *Manageability.* It is far from clear that implementing large-scale ubiquitous-computing scenarios involving potentially millions of smart, adaptive devices, is simply a question of scaling up existing toy examples. Will these services and applications still be able to meet their original requirements, even with a massive increase in the number of tiny interacting objects? And, above all, how will we be able to understand and control such a highly dynamic world involving such large numbers of individual objects?
- *Predictability.* Today's technical infrastructures, such as the phone system, television, and electricity, are relatively easy to use, even for people with no special qualifications. This also entails the ability to detect malfunctions: for example, if you lift a telephone receiver and do not hear a dial tone, it is immediately evident that the phone (either the handset or the landline) is not working properly. However, this type of predictability of system behavior can no longer be taken for granted in an ambient-intelligence landscape, as systems are expected to function without users noticing their presence. This will make fault detection and diagnosis fundamentally difficult, especially for the layman (Estrin *et al.* 2002). Additionally, users might continue to rely on a failed service (e.g., an automated backup service or the self-diagnostics of a smart product) without noticing, thus increasing the damage done until the problem is finally discovered.
- *Dependability.* Incorporating computing and communication technology into everyday artifacts requires ever-decreasing form factors and minimal energy consumption. This makes it difficult to use hardware redundancy in such systems, even though their envisioned unobtrusive and ubiquitous use implies much harsher surroundings than, say, an everyday indoor environment. This calls for alternative concepts and mechanisms in order to overcome service interruptions and device failures, such as an explicit diversification of system functions. Such a diversification can be achieved by providing fully independent ways of carrying out the same task, preferably based on separate sets of system resources wherever feasible. A communications connection, for instance, can be diversified if the system provides different communications mechanisms in parallel, such as GSM, Bluetooth, and wireless LAN.

The power outages that affected not only large parts of the USA and Canada but also Italy and some other countries in 2003 have demonstrated our dependence on existing technical infrastructures, in this case the power grid. With the constant goal of saving costs, any industry-built pervasive-computing infrastructure will run a high risk of forgoing safety for the sake of efficiency, resulting in brittle systems that will work only sporadically.

Delegation of Control

In order to minimize the need for human intervention in complex, highly dynamic environments, new concepts for delegating control are necessary – automatic processes should take care of routine tasks in a dependable manner, but also provide accounting mechanisms for monitoring complex control flows. Control and accounting mechanisms are important tools for determining who is in control of an autonomous system, and who is responsible if something goes wrong. At the same time, however, the autonomy of artifacts is also limited by their reliance on the technical infrastructure:

- *Content Control.* If smart objects provide information about themselves, this raises the question of who guarantees the objectivity and accuracy of the statements made. For example, smart products might be used to tie customers more closely to traders by recommending they purchase other goods produced by that same trader. In a certain sense, smart objects are becoming media representing a particular "ideology" (e.g., that of the product's manufacturer, or the politically motivated opinion of a consumer protection organization). Could a consumer protection institute use its own electronic directory to map a smart product label onto information other than that which the producer intended (for example, to warn of allergies to ingredients)? And 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.* It is similarly conceivable that automobiles or other products, as components of a pervasive-computing network, would no longer feel completely "loyal" to their owners, but would instead enforce the guidelines of insurance companies, manufacturers, or the judiciary. For example, a smart car might refuse to open the door for its driver because he or she has stopped in a no-parking zone. But when should an intelligent device obey human orders, and when should it follow its own "convictions"? While such a no-parking system might be desirable for congested cities, some kind of manual override mechanism would obviously be needed for emergency situations, e.g., when rushing a seriously injured person to hospital (and trying to park in front of it). Even if a system were designed to only make suggestions, it would still find itself treading a fine line between inspiration and frustration, between obliging helpfulness and pig-headed patronization (Satyanarayanan 2001).
- *Accountability.* If autonomous objects such as the previously mentioned smart doll start taking decisions on their own (e.g., buying new clothes), legal guidelines need to be drawn up in order to resolve who is ultimately responsible for these business transactions. Smart assistants might order unwanted plane tickets, smart fridges excessive amounts of food – in both cases the automated system might have performed according to specification, though neither the original programmers nor the layman user would be able to understand its reasoning. Providing the user with a detailed explanation of completed transactions is only part of the solution, especially when monetary damages are involved. It may look appealing to simply shift the responsibility and liability onto the end user by changing the license agreements of smart objects accordingly in the small print, but it is questionable whether such a procedure would prove tenable if taken to court.

Similar discussions involving the questions of accountability and content control are already taking place in the context of the World Wide Web. For example, questions regarding the right to possess and use certain prestigious domain names (Bitlaw 2000; Domain Name Handbook 2003) can be compared to the issue of content control (i.e., who is allowed to resolve a certain URL stored in the RFID tag of a product), while national laws trying to control digital copyright as well as freedom of speech (EFF 2003) might already set standards regarding the future "freedom" of smart devices to obey their owners.

Social Compatibility

Another fundamental challenge for pervasive-computing systems is their social compatibility. If we, as humans, want to be capable of participating in highly dynamic systems, their parameters will have to be adjusted accordingly. System behavior relating to particular aspects should retain a certain transparency and inertia, allowing humans to detect and adjust to changes. On the other hand, it should also be taken into account that an all-encompassing ambient-intelligence landscape must also meet the needs and requirements of as broad a section of society as possible, especially if participation is practically mandatory.

- *Transparency.* With the pay-per-use model discussed above, perusing an ambient-intelligence environment might incur a large number of micropayments, e.g., for bus or theater seats we have sat on, pages of books or newspapers we have read, or clothing we have worn. Irrespective of technical feasibility, this prompts the question of how we could keep track of the resulting number of short-term contracts and the countless associated micropayments, let alone retrospectively check the legitimacy of these transactions. Not only would it be extremely tedious and unrealistic to manually check thousands of transactions, it is also questionable to what extent inappropriate items could be identified and rejected, and to what extent legitimate payments could be unambiguously and indisputably allocated to the responsible party. Dynamic insurance rates that vary according to the style of driving (as discussed in an earlier section) constitute another example of a potential loss of transparency, especially if the underlying assessment methods changed dynamically with no warning, or if they were unknown or too complicated to be understood by the user.

- *Knowledge Sustainability.* Most information in our everyday life today remains valid for an extended period of time, e.g. food prices in our favorite supermarket, or prices for public transport. It is this inertia of information that permits us to use acquired knowledge and prior experiences to cope with future situations and tasks. 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 next. Such a loss or accelerated devaluation of long-term experiences could, in the long term, contribute to an increased uncertainty and lack of direction for people in society.
- *Fairness.* Detailed cross-marketing based on pervasive computing promises tailor-made offers that virtually eliminate unwanted advertisements. However, a specific offer may be withheld from a particular consumer for one of two reasons: either the offer was not worth the consumer, or the consumer was not worth the offer. 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 differences and reinforce existing inequalities" (Lyon 2001, p. 174). People not matching a certain 'desirable' profile might have to pay much higher prices, as they do not qualify for any of the existing discounts, which might in turn reinforce the non-matching patterns.
- *Universal Access.* The natural interfaces envisioned in ubiquitous-computing scenarios certainly have the opportunity to overcome many of today's accessibility problems, such as the small screens and keypads of modern mobile phones that often prevent elderly people from using them (Gonçalves 2001). Many projects in the field target elderly and physically disabled people in particular, for example with electronic "memory aids," reading aids and navigation systems (Makris 2001), which might pave the way for a universal design (Stephanidis 2001) that considers the needs of minorities and marginal groups early on in the design stage. Intelligent interfaces and the concept of ubiquitous information access are often seen as key developments for bridging the digital divide, where different sections of the population have different abilities to participate in the information society. However, having more information opportunities does not necessarily mean more justice or freedom, simply because the potential dependencies and opportunities for manipulation would be so numerous they could overwhelm individuals, making it even more difficult to assess the trustworthiness of the information's source. Information that was uncritical or sponsored by advertisers (and therefore one-sided) could become available free of charge, while independent, high-quality information would cost money, thus widening the digital divide even further. Since ubiquitous computing is not just about information itself, but is inherently linked to real-world objects, these new means of access and content control could easily lead to the digital divide becoming a real and perceivable rift in our everyday lives.

In history, the development of regulatory, social, and ethical standards tends to lag considerably behind the rapid proliferation of pioneering technological inventions, as was the case with the invention of the assembly line and mass production at the beginning of the 20th century, and with the appearance of the global Internet in the 1980s, for example. In an emerging future of pervasive-computing systems, one exciting question is whether we will be aware of the impending pitfalls and tackle them in an early (design) phase, where we still have the means to shape the envisaged systems according to fundamental social and ethical requirements, or if there is a need for yet another social revolution that subsequently brings about necessary adaptations by force.

Acceptance

The fundamental paradigm of pervasive and ubiquitous computing, namely that computers disappear from the user's consciousness and recede into the background, is sometimes seen as an attempt to have technology infiltrate everyday life unnoticed by the general public in order to circumvent any possible social resistance (Araya 1995, p. 236). Yet beyond any perceived sinister motives (which might be easy enough to counter), a widespread public acceptance of ubiquitous computing also rests on issues of an almost philosophical nature, such as the fundamental nature of smart objects or our changing relationship with our environment.

- *Feasibility and Credibility.* Many philosophers and social scientists identify a prevailing self-confident and technophile attitude among scientists in the field of pervasive computing (Adamowsky 2003), where the non-critical anticipation of future technological developments almost attains the characteristics of a metaphysical prophecy. Others doubt the credibility of the envisioned scenarios, e.g., when ubiquitous computing is said to simplify our lives, help us save time, and relieve us of laborious tasks. While this assertion has been constantly repeated throughout the twentieth century by the consumer goods industry, adding "smart machines" everywhere will not help to overcome the existing pattern of hurry, rush, stress, and separation from other people, but will only increase their efficiency (Winner 1999). Such criticism may build up and induce a serious credibility gap, reducing the acceptability of ubiquitous computing technologies.
- *Artifact Autonomy.* Networked everyday objects embedded in an ambient-intelligence landscape lose part of their autonomy and, with this, exhibit an increased dependence on the infrastructure. For users, this reduces the "object constancy" of the objects that surround them, as the example of electronic books made from smart paper shows: reading such a book may presuppose a regular connection to a server (license server, accounts server, etc.). Because of this, it appears to be more error-prone and less autonomous than a "normal" book, which can always be read, whereas the electronic one can only be read if the infrastructure is functioning.
- *Impact on Health and Environment.* It is hard to predict the impact that a large-scale use of ubiquitous computing and communication technology would have on our environment in terms of raw material consumption, energy consumption, and disposal. For example, if all supermarket goods were equipped with smart labels in the future, billions of these tiny and individually quite harmless chips would end up in the household garbage. On the other hand, the remote identification capability provided by smart labels would enable information on products to be made available throughout their entire lives, permitting the different materials in waste products to be efficiently identified and separated. It is also not yet fully understood whether, and to what extent, electromagnetic radiation (e.g., produced by wirelessly communicating smart objects) could affect our physical health. A vision involving myriads of everyday objects and wearable "information appliances" that communicate wirelessly with each other thus gives due cause for concern, as its potential adverse environmental effects could permanently influence the lives of future generations (Hilty *et al.* 2003).
- *The Relationship between Man and the World.* From a philosophical point of view, the vision of ubiquitous computing fundamentally changes the environment in which we live: "By this weaving of extensions of ourselves into the surroundings, significant parts of the environment lose important aspects of their otherness and the environment as a whole tends to become more and more a subservient 'artifact'. This artifact, which the world immediately surrounding us becomes, is almost entirely 'us' rather than 'other'. In this sense, the surrounding world has almost disappeared." (Araya 1995). Similarly, Adamowsky stipulates that our inability to handle the physical world in a flexible enough way will force us to replace it by digital surrogates – equivalents of particular aspects of the real world in the digital world, implemented in the form of models, simulations, and virtual counterparts – which will ultimately lead to a transformation, dislocation, substitution, and the loss of fundamental properties relating to the world (Adamowsky 2000).

Dryer *et al.* conducted two empirical studies to examine the theoretical relationships between system design for mobile computing, human behaviors, social attributions, and interaction outcome. In their conclusion (Dryer *et al.* 1999, p. 674), they express "doubt that our inevitable future is to become a machinelike collective society. How devices are used is not determined by their creators alone. Individuals influence how devices are used, and humans can be tenaciously social creatures." They conclude "Given the importance of social relationships in our lives, we may adopt only those devices that support, rather than inhibit, such relationships." With the substantial amount of skepticism related to technology, such findings seem to counterbalance the immediate threat that a thoroughly computerized future appears to hold. However, apart from personal prejudices, the wide range of social consequences that pervasive computing may have will certainly need to be addressed in future systems and debates. These challenges are of fundamental importance and may ultimately even have a decisive influence on the large-scale acceptance of ubiquitous-computing technologies and environments.

BRAVE NEW WORLD?

"Everything will be connected to everything else," but "no one has any idea what all those connections will mean" (Lucky 1999, p.19). This criticism can be taken as a perceived lack of focus when it comes to pervasive-computing applications, but also as a deficiency in terms of understanding the consequences of deploying pervasive-computing systems in the real world: how will we use "smart things" in our everyday lives? When should we switch them on or off? What should smart things be permitted to hear, see, and feel? And whom should they be allowed tell about it? Whether these consequences concern the protection of personal data, the implications for the macro-economy, or social acceptance – developers of pervasive-computing systems can profit greatly from a careful evaluation of the consequences of such technology within the framework of established concepts from the fields of sociology, economics, and jurisprudence.

Although predicting the future is difficult, if not impossible, the above discussion allows us to guess at a few of the possible implications of wide-scale use of ubiquitous-computing technology: new business models will increase profits, possibly at the expense of safety margins; the balance of political and economic power could shift; economic developments will accelerate and initiate long-term changes in our social values and motives; personal borders could be violated by new surveillance and search technology; and, not least, there is the danger that we will lose confidence in our environment, thus fundamentally and unfavorably changing our attitude towards the world that surrounds us. The intention of this article was to throw light on the interdisciplinary field of pervasive computing, in order better to understand how far these visions can and should influence our everyday lives. By identifying and addressing the great challenges of technical and social change, as well as their environmental sustainability, it may be possible to steer this development in a direction that has more in common with Weiser's optimistic vision of the 21st century than with the depressing mix of consumer terror and police state conjured up by Steven Spielberg in his movie "Minority Report" (Dick 1956).

ACKNOWLEDGMENTS

This work has been partially funded by the Gottlieb Daimler- and Karl Benz-Foundation, Ladenburg, Germany, as part of the interdisciplinary research project "Living in a Smart Environment – Implications of Ubiquitous Computing". The authors also wish to thank Jochen Jagob, TU Darmstadt, Germany, for fruitful discussions on the economic issues.

REFERENCES

Abowd GD, Brumitt B, Shafer S (eds). 2001. Ubicomp 2001: Ubiquitous Computing. Proceedings Series: Lecture Notes in Computer Science, Vol. 2201, Springer-Verlag, Berlin Heidelberg New York

Accenture Technology Labs. 2001. Silent Commerce. Available at www.accenture.com/xd/xd.asp ?it=enweb&xd=services/technology/vision/silent_commerce.xml

Adamowsky N. 2000. Kulturelle Relevanz. Ladenburger Diskurs „Ubiquitous Computing". Available at www.vs.inf.ethz.ch/events/slides/adamowldbg.pdf

Adamowsky N. 2003. Smarte Götter und magische Maschinen – zur Virulenz vormoderner Argumentationsmuster in Ubiquitous-Computing-Visionen. In: Mattern F (Mattern 2003), pp 231–247

Ahola J. 2001. Ambient Intelligence. ERCIM News, No 47. Available at www.ercim.org/publication/ Ercim_News/enw47/

Akerlof G. 1970. The Market for Lemons: Qualitative Uncertainty and the Market Mechanism. The Quarterly Journal of Economics, 84(3), pp 488–500

Araya AA. 1995. Questioning Ubiquitous Computing. In: Proceedings of the 1995 ACM 23rd Annual Conference on Computer Science. ACM Press. Available at doi.acm.org/10.1145/259526.259560

Bitlaw. 2000. Some well publicized examples of domain names disputes. See www.bitlaw.com/internet/domain.html#disputes

Cochrane P. 2000. Head to Head. Sovereign Magazine, pp 56–57, Spring 2000

Coroama V, Röthenbacher F. 2003. The Chatty Environment – Providing Everyday Independence to the Visually Impaired. In: 2nd Int. Workshop on Ubiquitous Computing for Pervasive Healthcare Applications (UbiHealth2003). Available at www.vs.inf.ethz.ch/publ/papers/ubicomp2003-hc.pdf

Council of Europe. 2001. Convention on Cybercrime. ETS No 185. Available at conventions.coe.int/Treaty/en/Treaties/Html/185.htm

Dick PK. 1956. Minority Report. Fantastic Universe, January 1956

Dryer DC, Eisbach C, Ark WS. 1999. At what cost pervasive? A social computing view of mobile computing systems. IBM Systems Journal, 38(4):652–676

Domain Name Handbook. 2003. Domain Dispute Index. See www.domainhandbook.com/dd.html

Electronic Frontier Foundation EFF. 2003. Blue Ribbon Campaign for free speech online. See www.eff.org/br/

Estrin D, Culler D, Pister K, Sukhatme G. 2002. Connecting the Physical World with Pervasive Networks. IEEE Pervasive Computing – Mobile and Ubiquitous Systems, 1(1):59–69

Ferguson GT. 2002. Have Your Objects Call My Objects. Harvard Business Review, 80(6):138–144

Finkenzeller K. 2003. RFID-Handbook, 2nd edition. John Wiley & Sons, Chichester

Fleisch E. 2002. Von der Vernetzung von Unternehmen zur Vernutzung von Dingen. In: Schögel M, Tomczak T, Belz C (eds). 2002. Roadm@p to E-Business – Wie Unternehmen das Internet erfolgreich nutzen. Thexis, St. Gallen, pp 124–135

Gonçalves DJ. 2001. Ubiquitous Computing and AI Towards an Inclusive Society. In: Heller *et al.* 2001, pp 37–40

Heller R, Jorge J, Guedj R (ed). 2001. Proceedings of the 2001 EC/NSF Workshop on Universal Accessibility of Ubiquitous Computing: Providing for the Elderly, Alcácer do Sal, Portugal, May 2001, ACM Press. See also virtual.inesc.pt/wuauc01/

Hilty L, Behrendt S, Binswanger M, Bruinink A, Erdmann L, Fröhlich J, Köhler A, Kuster N, Som C, Würtenberger F. 2003. Das Vorsorgeprinzip in der Informationsgesellschaft. Auswirkungen des Pervasive Computing auf Gesundheit und Umwelt. Studie des Zentrums für Technologiefolgen-Abschätzung TA-SWISS, TA 46/2003

Joshi YV. 2000. Information Visibility and its Effect on Supply Chain Dynamics. Master's Thesis, MIT

Lee HL, Padmanabhan V, Whang S. 1997. The Bullwhip Effect in Supply Chains. MIT Sloan Management Review, 38(3):93–102

Lessig L. 1999. Code and Other Laws of Cyberspace. Basic Books, New York NY

Lucky R. 1999. Everything will be connected to everything else. Connections. IEEE Spectrum, March 1999, p 19. Available at www.argreenhouse.com/papers/rlucky/spectrum/connect.shtml

Lyon D. 2001. Facing the Future: Seeking Ethics for Everyday Surveillance. Ethics and Information Technology, 3(3):171–180

Maeder T. 2002. What Barbie Wants, Barbie Gets. Wired Magazine, 10(1):6

Makris P. 2001. Accessibility of Ubiquitous Computing: Providing for the Elderly. In: Heller *et al.* 2001

Marx GT. 2001. Murky Conceptual Waters: The Public and the Private. Ethics and Information Technology, 3(3):157–169

Mattern F (ed). 2003. Total Vernetzt: Szenarien einer informatisierten Welt. Springer-Verlag, Berlin Heidelberg

Mattern F. 2004. Ubiquitous Computing: Scenarios for an informatized world. In: Zerdick *et al.* 2004

Mayo RN. 2001. The Factoids Project. Available at www.research.compaq.com/wrl/techreports/abstracts/TN-60.html

Moore GE. 1965. Cramming more components onto integrated circuits. Electronics, 38:114–117

Nagel KS, Kidd CD, O'Connell T, Day A, Abowd GD. 2001. Family Intercom: Developing a Context-Aware Audio Communication System. In: Abowd *et al.* 2001, pp 176–183

O'Harrow Jr R. 2001. Prozac Maker Reveals Patient E-Mail Addresses. The Washington Post, July 4, 2001

Reuters. 2003. This DVD Will Self-Destruct. Wired News, May 16. Available at www.wired.com/news/technology/0,1282,58883,00.html

Rhodes B. 1997. The Wearable Remembrance Agent: A System for Augmented Memory. Personal Technologies Journal. Special Issue on Wearable Computing, 1:218–224

Rhodes B, Minar N, Weaver J. 1999. Wearable Computing Meets Ubiquitous Computing – Reaping the Best of Both Worlds. In: Proceedings of the Third International Symposium on Wearable Computers (ISWC '99), San Francisco CA, pp 141–149

Rotenberg M. 2001. Testimony and Statement for the Record. Hearing on Privacy in the Commercial World before the Subcommittee on Commerce, Trade, and Consumer Protection, U.S. House of Representatives, March 2001. Available at www.epic.org/privacy/testimony_0301.html

Satyanarayanan M. 2001. Pervasive Computing: Vision and Challenges. IEEE Personal Communications, 8(4):10–17

Siegele L. 2002a. Tante Emma lebt. Die Zeit, (42):27

Siegele L. 2002b. How about now? A survey of the real-time economy. The Economist, 362(8257):3–18

Skiera B, Spann M. 2002. Preisdifferenzierung im Internet. In: Schögel M, Tomczak T, Belz C (eds). 2002. Roadm@p to E-Business – Wie Unternehmen das Internet erfolgreich nutzen. Thexis, St. Gallen, pp 270–284

Stephanidis C. 2001. Towards Universal Access in the Information Society. In: Heller *et al.* 2001

Talbott S. 2000. The Trouble With Ubiquitous Technology Pushers, or: Why We'd Be Better Off without the MIT Media Lab. NetFuture: Technology and Human Responsibility. Available at www.netfuture.org/2000/Jan0600_100.html#3

Thackara J. 2001. The design challenge of pervasive computing. Interactions, 8(3):46–52. Available at doi.acm.org/10.1145/369825.369832

USA. 1998. The Declaration of Independence and the Constitution of the United States. Bantam Books, New York

USA Today. 2000. Amazon May Spell End for 'Dynamic' Pricing. Associated Press. September 29, 2000

Warren S, Brandeis L. 1890. The Right to Privacy. Harvard Law Review, 4(1):193–220

Weiser M. 1991. The Computer for the 21st Century. Scientific American, 265(3):66–75

Weiser M. 1993. Ubiquitous Computing. IEEE Computer, 26(10):71–72

Westerlund B, Lindquist S, Sundblad Y. 2001. Cooperative Design of Communication Support for and with Families in Stockholm. Available at interliving.kth.se/papers.html

Westin AF. 1967. Privacy and Freedom. Atheneum, New York NY

Winner L. 1999. The Voluntary Complexity Movement. NetFuture: Technology and Human Responsibility, September 1999. Available at www.netfuture.org/1999/Sep1499_94.html#3

Zerdick A, Picot A, Schrape K, Burgelman J-C, Silverstone R (eds). 2004. E-Merging Media, pp 155–174, Springer-Verlag, Berlin Heidelberg New York