As we may live – Real-world implications of ubiquitous computing

Marc Langheinrich, Vlad Coroama, Jürgen Bohn, and Michael Rohs

Distributed Systems Group
Institute of Information Systems
Swiss Federal Institute of Technology, ETH Zurich
8092 Zurich, Switzerland
www.inf.ethz.ch/˜{langhein,coroama,bohn,rohs}

Abstract. The young field of ubiquitous computing is steadily making progress and gaining attention in both academia and industry. While new gadgets and smart home appliances cannot appear fast enough for many technologists, such rapid introductions of new technologies often come with unexpected side-effects. Due to the unique scope of ubiquitous computing as a tool for our everyday life, these side-effects might have serious implications for the way we live in the future. This paper explores a number of effects that a large-scale deployment of ubiquitous computing technology in the real world may have. Its intention is to raise awareness for a technical design that takes the concerns of the average citizen into account, as well as to serve as point of departure for further interdisciplinary research in the field.

1 A brave new world

More than a decade ago, Xerox PARC researcher Mark Weiser coined the term and defined the field in his seminal work “The computer for the 21st century” [46]. After Vannevar Bush’s “As we may think” [7] set the tone for “a new relationship between thinking man and the sum of our knowledge,” Weiser’s ideas shifted the focus from Bush’s virtual world towards the relationship between our *lives* and the sum of our *technology* [13].

While such ideas seemed slightly utopian then, the availability of single-chip wireless communication solutions (e.g., Bluetooth [5]), Java-enabled Smartcards [17], or radio-frequency-based identification systems (RFID-Tags) with a form factor of only a few micrometers [16] has allowed researchers around the world to start putting their ideas to work and creating a large variety of prototypes with ease [37].

Commercial interest has also been picking up. After the lackluster development of mobile commerce in many markets, industry is looking for better ways to turn its investments in telecommunication sectors into profit, and “smart devices” that provide “information at the user’s fingertips” might very well be the key to increase consumer acceptance.

Consequently, research funding both in academia and industry is plenty and a wealth of prototypes and field tests appear in all parts of the world, heralding a new age of “invisible computing” that Mark Weiser and his peers envisioned more than ten years ago.

But while scientists and engineers are busy imagining new application domains for the plethora of “cool stuff” that makes up their prototypes, the details of these developments go often unnoticed by the public at large.

Yet for all its “geeky” touch, ubiquitous computing will very probably be far more consequential for our daily life, social values, and core beliefs, than the Internet with all its discussions about unsolicited advertising, cybercrime [31], and child pornography could ever be. With its large applicability across both public and private, personal and business domains, developments in ubiquitous computing will potentially affect all of our life, all of the time. And if Mark Weiser’s vision gets properly executed, we won’t even notice!

It is the aim of this paper to stimulate discussion in the interdisciplinary borderland surrounding ubiquitous computing and venturing into non-technical fields like sociology, economics, or law. By taking into account the research available in these related disciplines, and directly connecting professionals in these fields with the research efforts underway in this fast-moving area of “invisible computing,” synergies may be found that help us channel our development efforts into the right directions more closely resembling Weiser’s vision of the 21st century than Orwell’s 1984.

The remainder of this paper is organized as follows: Section 2 begins with looking at implications of large-scale deployment of ubiquitous computing with respect to personal privacy. Section 3 broadens the scope and describes economic implications of a world full of “smart things,” including both business management and macro-economic aspects. Section 4 will then examine how we are to depend on a thoroughly interconnected environment, while section 5 sums up popular critique of ubiquitous computing and examines it in light of our findings in the three previous sections. Closing arguments and final assessments can be found in section 6.

2 Left to your own devices – Personal privacy in ubiquitous computing

As the field of ubiquitous computing matures, more and more of the key issues start shifting away from mere technical problems to those that have a fundamentally *social* background: How are we to use those smart devices in our daily routine? When should they be turned on and off? What should they be allowed to see, feel, or hear? And whom should they tell about it?

Among such questions, privacy is probably the most prominent concern when it comes to judging the effects of a widespread deployment of ubiquitous computing. This is certainly due to the already imminent threat to privacy caused by the ever growing use of distributed commercial databases that record large parts of our daily electronic transactions. By virtue of its very definitions, ubiquitous computing has now the potential to create an even more invisible and comprehensive surveillance network covering an unprecedented share of our public and private life. Consequently, much has been written about privacy in light of automated data processing [6, 11, 10], though less so in the context of ubiquitous computing [4, 21].

The following sections try to add a more differentiated view on the impact of ubiquitous computing on personal privacy by first examining *why* personal privacy is desirable, describing *when* we feel that it has been violated, and then assessing *how* ubiquitous computing affects all that.

2.1 A private affair – Motivating personal privacy

Along with articles covering privacy aspects, a range of definitions for what actually constitutes privacy are given, the most prominent probably being judge Brandeis' "The right to be left alone" [45] and Alan Westin's "The claim of individuals... to determine for themselves when, how, and to what extent information about them is communicated to others" [50]. These definitions certainly help to illustrate that privacy not only has different goals in different contexts, but also that personal limits for privacy differ according to factors such as geography (e.g., whether we are at home or in a public park), informational access rights (e.g., anti-mask laws in certain states/countries prohibit hiding ones face in public), or expectations and manners (e.g., expecting people not to openly stare at you in public) [26]. Depending on any of these dimensions, individuals can both expect a reasonable level of protection from the prying eyes and ears of their fellow citizens, or be required to disclose certain parts of their own information when necessary by law or custom.

A valuable avenue for exploration when trying to assess the implications of new technology on something as old as the concept of privacy, might be to look at the *motivations* behind privacy, as it is far from undisputed that societies in fact need the level of privacy protection that its most ardent proponents would like to have. Scott McNealy's infamous "You have no privacy anyway, get over it" [39] and Peter Cochrane's "All this secrecy is making life harder, more expensive, dangerous and less serendipitous" [8] indicate a growing backlash among those tired of hearing the constant warnings coming from privacy advocates.

Privacy is often seen as a fundamental requirement for any modern democracy [35]. Only if people can freely choose according to their interests and believes, without fear of repression from their fellow citizens, the necessary plurality of ideas and attitudes can grow that prevent bringing the general public into line by charismatic leaders. Harvard law professor Lawrence Lessig [24] takes this requirement a step further and differentiates between a number of motivations for privacy protection in our present-day laws and norms:

- **Privacy as empowerment:** Seeing privacy mainly as informational privacy, its aim is to give people the power to control the dissemination and spread of information about themselves. A recent legal discussion surrounding this motivation revolves around the question whether personal information should be seen as a private property (which would entail the rights to sell all or parts of it as the owner sees fit) or as intellectual property (which would entitle the owner to certain unalienable rights, preventing him for example to sell the rights to his name to anybody).
- **Privacy as utility:** From the data subject's point of view, privacy can be seen as a utility providing more or less effective protection from nuisances such as unsolicited calls or emails. This view probably best follows Brandeis' "The right to be left alone" definition of privacy, where the focus is on reducing the amount of disturbance for the individual.

- **Privacy as dignity:** Dignity can be described as “the presence of poise and self-respect in one’s deportment to a degree that inspires respect.” [32] This not only entails being free from unsubstantiated suspicions (for example when being the target of a wire tap, where the intrusion is usually not directly perceived as a disturbance), but rather focuses on the *balance* in information available between two people: analogous to having a conversation with a fully dressed person while being naked oneself, any relationship where there is a considerable information imbalance will make it much more difficult for those with less information about the other to keep one’s poise.
- **Privacy as constraint of power:** Privacy laws and moral norms to that extend can also be seen as a tool for keeping checks and balances on a ruling elite’s powers. By limiting information gathering of a certain type, crimes or moral norms pertaining to that type of information cannot be effectively enforced. As Stunz [40] puts it: “Just as a law banning the use of contraceptives would tend to encourage bedroom searches, so also would a ban on bedroom searches tend to discourage laws prohibiting contraceptives.”
- **Privacy as by-product of imperfect surveillance tools:** While law enforcement in many democratic countries can in principle search any private premises, listen in to any private phone call, and open any number of private letters, given a proper search warrant, their actual ability to do so is often quite limited: searches and surveillance takes both time and money, so officers usually try to make sure they spend their efforts on some reasonably suspicious target. The larger, unsuspicious looking general public must thus rarely consider themselves the target of such a surveillance or search, simply because the effort would hardly be worth it. The resulting level of privacy they consequently enjoy stems inasmuch from the required court-ordered warrant, as from the imperfection of the employed search and surveillance tools.

Depending on what kind of motivation one assumes for preserving privacy, ubiquitous computing can become the driving factor of changing the reach and impact of privacy protection as it exists today, and create substantially different social landscapes in the future. It can do so because ubiquitous computing influences two important design parameters relating to privacy: the ability to *monitor* and the ability to *search* [24].

2.2 Lookin’ good - Ubiquitous computing and surveillance

Monitoring people and their actions and habits is a human trait as old as humanity itself. In the “good old days”, such monitoring would constantly be done within small villages and settlements by our close social peers, who would immediately notice anything out of the ordinary and disseminate it in society. It was this close monitoring that often enough drove people into the big cities, where the sheer number of citizens and their constant mobility effectively put an end to the watchful eyes of the neighbors. Yet with the advent of automated information processing, machines took over the role of the watchers and began to store more and more of our daily routines, not only when they happened to be “out of the ordinary.” With ubiquitous computing, monitoring capabilities can obviously be extended far beyond credit-card records, calling logs, and news postings. Not only will the *spatial* scope of such monitoring activities be significantly extended with ubiquitous computing but also their *temporal* coverage will vastly increase: starting from pre-natal-diagnostics data stored on the baby's health-id-card, to activity feeds in kindergarten and schools, to workplace monitoring and senior citizen 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." [26]. He goes on to define four such border crossings that form the basis for perceived privacy violation:

- **Natural borders:** Physical limitations of observations, such as walls and doors, clothing, darkness, but also sealed letters, telephone calls. Even facial expressions can form a natural border against the true feelings of a person.
- **Social borders:** Expectations about confidentiality for members of certain social roles, such as family members, doctors, or lawyers. This also includes expectations that your colleagues will not read personal fax messages addressed to you, or material that you left lying around the photocopy machine.
- **Spatial or temporal borders:** The usual expectations of people that parts of their life, both in time and social space, can remain separated from each other. This would include a wild adolescent time that should not interfere with today's life as a father of four, or different social groups, such as your work colleagues and friends in your favorite bar.
- **Borders due to ephermal or transitory effects:** This describes what is best known as a "fleeting moment," an unreflected utterance or action that we hope gets forgotten soon, or old pictures and letters that we put out in our trash. Seeing audio or video recordings of such events later, or observing someone sifting through our trash, will violate our expectations of being able to have information simply pass away unnoticed or forgotten.

Putting ubiquitous 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* [27, 33], allowing its wearer to constantly record events of her daily life in a lifetime multimedia diary. While at first sight such a technology promises great help for those of us who tend to forget a lot of 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 as a high-quality audio and video feed if my conversation partner would give others a peek into her multimedia diary. Even if this information would never get disclosed to others, just the thought of dealing with people who have a perfect memory and in theory would *never* forget anything, will probably have a sizable effect on interpersonal relationships.

The problem of *spatial and temporal borders* on the other hand is well known from the area of consumer profiles. Profiles are often enough the focus of public concerns, but so far social and legal attitudes have been relatively relaxed about them. Consumer acceptance is also much higher than the often negative news coverage might indicate, mostly because their harm is often perceived as being small (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 widespread as the basis for a ubiquitous computing infrastructure will only intensify such problems. Besides the obvious risk of data spills [18], profiles also threatens universal equality, a concept central to many constitutions, basic laws, and human rights, where "all men are created equal." [43]. Even though a thoroughly customized future (using ubiquitous computing) where I only get the information that is relevant to my (very comprehensive) 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, constitutes a severe privacy violation for many people.

Applying ubiquitous computing technology in areas with primarily *social borders* – for example where a close social group interacts only among themselves, such as families [29, 49] or co-workers – might seemingly alleviate some of the above concerns. Most participants share already 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 also increases the risk for unwanted social border crossings by accidentally allowing Mum to read a message you left for your sister, or a visiting friend to appear in the house activity log even though you told grandma you would spend the weekend alone.

Natural borders, then, might be easiest to respect when designing ubiquitous 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 users belt, or within her shirt) as a turnkey solution for privacy conscious technologists [34]. Border crossings, however, are not only about *who* does something, but *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 it might find there might easily startle (and trouble!) us, once such systems would start anticipating our future actions and reactions. The feeling of having someone (or something) constantly peeking 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 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 laws against cybercrime [31] do this for computer log files).

2.3 Don't ask, don't tell – The power of search

All these examples serve to show that ubiquitous computing systems, even when installed for the greater good and with the best of intentions, will run a high chance of involuntarily threatening our personal borders that set apart private from public, simply because their monitoring capabilities will facilitate more of the border crossings described above. Whether or not such crossings ultimately occur, given the opportunities created, will to a large extend also depend on the type of *searching* capabilities that such ubiquitous computing systems might offer.

Search technology is traditionally not a particular focus of ubiquitous computing, mainly since its core methods are more likely to be developed in the fields of information retrieval or databases. However, what *will* become relevant in ubiquitous computing is how the chosen architectures will support such search techniques. Chances are high that such technology will be a basic building block of future ubiquitous computing systems, as most of the envisioned applications in the fields of *context-awareness* and *memory augmentation* require just these capabilities. An automated diary collecting 24/7 audio and video-feeds will not be of much use unless being combined with a powerful search and retrieval technology that lets us 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 envisioned "smart" system, which must make best use of a large variety of different sensor input to come to decisions that make it appear as if it would *understand* what was happening around us.

Having thus both monitoring and search capabilities at the very core of their architecture, ubiquitous computing system will very likely provide their developers, owners and regulators with a significant tool to drive the future development of privacy concepts in society. Depending on the actual systems that receive large-scale deployment, some of the motivating aspects of privacy as discussed in section 2.1 might become more or less prominent, thus influencing corresponding legal and social norms.

For example, imagine law enforcement having a low-cost ability to search a large number of homes without effort in short time, for example by having all home automation manufacturers build in hooks into their software that would allow police to register certain behavioral patterns and let motion, audio and video sensors report in when they detect a suspicious match. The temptation to try one's luck in order to find a certain suspect might very well lure policymakers, judges and police into giving up today's relatively cumbersome privacy laws, marking privacy as it exists today as a simple residue of inefficient tools that can be abandoned in favor of national security. By motivating privacy instead as a simple *utility* with a bit of *dignity* thrown in, these searches could still be considered privacy-friendly as they would neither inconvene those subject to such a search, nor would they report any personal actions that would not fit the registered suspicious behavior.

Examples for consequences in ubiquitous systems design then, given the above findings, are numerous. They could include commendations to use sense-enhancing technologies only sparsely in ubiquitous computing, and only in limited, well-defined environments (e.g., emergency room, aircraft hangars). Communication concepts could be evaluated according to the existent social borders of all participants, in order to prevent unwanted data spills. Searching capabilities that allow spatial and temporal border crossings would need to be questioned, and the concept of ephemeral, transitory effects be re-introduced into ubiquitous computing architectures, allowing for example that information slowly decays over time.

What is important to realize is that technical concepts alone often create only a superficial understanding when it comes to real-world implications of ubiquitous computing. This might become again apparent in the next section, where we go on to explore another set of possible real-world implications, those of the economy. Not only are economic issues a source of interesting application domains for ubiquitous computing, but they also have the potential to become a major driving force for the deployment of ubiquitous computing systems in the near future.

3 Just in time – Economic implications

Ubiquitous computing techniques enable us to model in computers the physical reality more closely than ever before. This growth takes place along two axes: the quantity axis (more and more parts of reality are modelled) and the time axis (the time between an event happening and being represented drops, up to becoming real-time).

In this section we examine how the economy could be affected by ubiquitous computing and this closer-than-ever-before reality representation. In particular, we will examine how companies can *improve existing business processes* by being able to track the location and status of goods both inside the company and along the supply chain, how they can create *completely new business models* with such information, and how these techniques may be used to *influence macro-economic effects* such as taxes.

3.1 Now or never – The instant economy

Even though the early hype surrounding e-commerce and b2b-transactions has ended, information technology and the internet have nevertheless significantly help reshape companies and their way of doing business. The concept of “real-time economy” or “now-economy” [38] expresses the fact that companies increasingly use sophisticated IT systems to gather extensive real-time information about the whereabouts of company entities and constantly monitoring their status, thus increasing the transparency of company assets and improving its reaction time to unforseen events.

Ubiquitous computing with its extensive monitoring and searching capabilities is only a natural extension of this trend, transforming “now-economies” into “instant-economies”, where location and status information of goods, people, and orders can be tracked instantaneously and in high precision.

One reason why such a transparency of assets can save money is the problem of inventory tracking. Without knowing which goods are where for how long in their warehouses, companies incur lost profits from sale, overstocked raw supplies, and diminishing asset value over time (e.g, from perishing food or outdated products).

Typically, such (periodically conducted) inventory assessments require a considerable amount of manual labor, involving a large number of employees and usually requiring a temporary suspension of operations. Not only is this costly (both in terms of salaries paid and profits lost), but it is highly error prone.

Using ubiquitous computing technology such as indoor location tracking or RFID tagging, such inventory tracking tasks could be completely automated, saving both the extra costs of human labor and suspended operation and increasing accuracy and ’freshness’ of the assessment at the same time.

In addition, companies can combine this up-to-date information about both assets and orders with other companies along their supply chain (e.g., producers of raw materials and wholesalers) and realize further saving by cancelling losses stemming from the “bullwhip-effect” [23]: even though consumer demand stays fairly constant, small variation in buying patterns are increasingly amplified at each link along the supply chain, resulting in either greatly exaggerating production (creating unwanted inventory) or sudden interruptions in supplies (requiring backordering).

A step further in instant economies represents the tracing of further product parameters, such as temperature, acceleration or pressure, using tiny wireless sensors embedded in their electronic product tags. Equipped with communication facilities, such “smart” goods are not only able to observe themselves, but also to independently communicate relevant parameters to the outside world. Aircraft turbines, for instance, are beginning to be equipped with sensors and thus permanently monitor themselves [38]. If they detect any unusual patterns during operation (which is typically inflight), an immediate order for the corresponding spare parts can be made to the destination airport, allowing the material to be already at hand when the airplane lands and thus minimizing turn-around time. Another example would be chemicals and food products that would constantly monitor their temperature during transport, and independently trigger an alarm or dynamically adjust their best-before date should they get spoiled in the process. This would not only improve consumer health but also avoid costly verification of products (e.g., chemicals) or their use as a raw material should their characteristics have changed (note that in this case resupplies could also be ordered more quickly).

With ubiquitous computing and information not only the current way business is done can be improved; completely new business models are imaginable, some with a strong impact on the way we will perceive economic reality in the future.

3.2 From super-markets to stock-markets – Just-in-time pricing

A perfectly competitive market (for short: perfect market) is defined in classic economic theory [19] as having three characteristics: homogenous goods; complete and correct information on both supply and demand side; and no time or space advantages for some of the market players over the others. Nowadays only stock-markets can be regarded as perfect markets; only here goods are interchangeable (one stock being as good as the other), all players are concentrated in one place at the same time, and all have complete information about prices being asked and offered. Advantages of a perfect market include peak trading of goods and optimum price for a majority of sellers and buyers. Such markets are characterized by a highly-dynamic price structure. There is a permanent bargaining to find the current market price; based on this price market players come to a decision for their next-moment moves, which in turn determine next moment’s price and so on.

This kind of permanent negotiation becomes more and more popular, as the example of online auction houses like *eBay* shows. It has advantages for both buyers (who hope to make a good bargain) and traders (who can get rid of stocks for a lower price or obtain better prices for demanded items than they would in their geographic neighborhood).

Ubiquitous computing techniques have the potential to transform many traditional, static marketplaces into highly-dynamic ones. One extreme example could be supermarkets. If all products in a super-market are able to sense their environment – e.g., other products around, the time of day, the day in the week, etc – and communicate with other products, with the shelves and maybe the cash register, then super-markets can become perfectly competitive markets where prices are determined dynamically, in correspondence to supply and demand. In short, *the super-market becomes a stock-market*.

Take for example a milk bottle equipped with sensing, computational and communication facilities. Bottles can communicate with each other and the market itself, so each one knows how many are on the shelf, if there is some supply in the warehouse and also the expiration dates of all other bottles. The shelf also knows all that plus other information influencing milk demand; like time of day, day of week, season or weather outside. And, of course, the history of milk buying over the past months or years.

Dependent on all these parameters, milk bottles will set their prices dynamically. For example, while approaching the expiration date, a bottle will decrease its price, so customers will be tempted to buy it and not grasp for another, fresher one. Same would happen on the third rainy day in sequence, when sales decrease and part of the stock risks to expire. On the other hand, when sales are exceeding expectations, bottles notice they become lonelier in the shelf with time passing by. If the warehouse is also empty, remaining bottles will steadily increase their prices, as long as people are buying. Thus, highly dynamic reaction to market facts is possible.

The technical implementation of such a system is non-trivial. To have the shelves displaying the price seems not a feasible solution. More likely, the goods themselves should have means to display their momentary price. One solution would be to enhance them with flexible displays or smart paper. Adding this to the already necessary tags, sensors, and communication moves this scenario further away in the future, at least for low-value products like super-market goods.

Of course, changing super-markets policy to this highly dynamic price finding will not necessarily be a financial success. People are customized to count on stable prices and do not want to spend their energies to continuously watch for bargains. They may be willing to bargain for high-value products like cars, or when buying a handheld computer over *eBay* instead in the store. Nevertheless, they may perceive it as a marginal value to have dynamic milk prices and could be scared about too many new things to pay attention to (like not paying $10 for a liter of milk because it happens to be Friday afternoon) and would like to go back to the *predicatbility* of annual holiday-sales. One answer could be enforcing by law or economic constraints that items may become less, but not more expensive. In any case, this issue deserves further research.

Last not least, the perfectly competitive markets envisioned above not only seem very profitable, but are also known to be rather innovation-unfriendly places. In very competitive markets, companies have usually only small margins of profits, and thus are often unable to heavily invest in research and new business ideas. The resulting general climate of defensive behavior might become a long-term drawback for a ubiquitous computing economy as well.

3.3 I owe you – Pay-per-use paradigm

Another business model that could potentially find more prevalence by means of ubiquitous computing is pay-per-use. Evolving digital rights management (DRM) systems represent the attempt of music and software companies to impose intellectual property rights. Through such systems it is possible to sell customers only restricted access to the data they are buying; the user might for instance listen to the CD she "bought" only after 6pm or just three times altogether.

Equipping various everyday objects with sensors and communication facilities could bring up a new dimension in pay-per-use businesses. Almost every product becomes suitable for pay-per-use instead of buying. Researchers at *Accenture*, for instance, have build a prototype of a pay-per-use chair capable of sensing the intervals different people used it [20].

Accenture Labs praise this model as being great for both seller and consumer: "Obviously, it's great for the buyer because they only pay for what they use." And further: "These embedded devices mean that almost anything can be pay-per-use". At first glance, this does not really sound like "great for the buyer". People are used to own things, not to lend them. Imagine an example similar to car leasing, where you used a sofa for more than the agreed upon 150 hours. As soon as the contracted hours are all used up, the furniture supplier (connected all the time with the sofa – after all it's still their piece of furniture!) would either come and pick it up from your living room or require you to renegotiate a follow-up agreement.

Another possible new business model are insurances with highly personalized and dynamic insurance rates. Criteria like the way you drive, letting others drive your car or not, their way of driving, often driving at night, or where you park the car, may be taken into account in order to calculate custom car insurance rates. Driving fast will not only increase your gas bill, the insurance company will also be glad to notice it.

Even if you theoretically still have the possibility to opt-out, at which price will that be? If people not willing to send their data to the insurer will have to pay three times the average insurance rates, most people will agree to do so. Money always has been a strong argument! That is why such business models have to be addressed cross-disciplinary; by also taking in consideration the possibility to limit through law some technical possible, but socially undesirable business models.

3.4 What you pay is what it took – Dynamic taxation and the economy

Government could also take advantage of the new technologies. By knowing the history of products, truly fine-granular taxation becomes feasible. Milk bottles, for instance, could determine ecologic taxes by themselves, being aware whether they've been transported by truck or by railway. Granulation can go even further: taxes could depend on the length of transportation, encouraging regional producers. Other product properties may depend on the history of product manufacturing. Milk would be rated as organic if the cow providing it has been fed exclusively by natural products and not been given antibiotics.

Fine control mechanisms can be imagined: to help a region ravaged by a calamity, government could reduce taxes for products shipped to this region. After taking the political decision to do so, the only technical issue is to propagate this information to the local software tax agents.

Not less important: the effects of these measures could be analyzed in real-time and accurate by deployment of ubiquitous computing. Two days after, say, increasing taxation for milk produced in a certain district, the effects will be reflected in decreased over-the-counter sales. Macroeconomic models trying to estimate such consequences would be consequently improved.

Analyzing data from a large number of sources may sometimes lead to unexpected results that couldn't have been found otherwise. *Economist's* "R-factor" is a well-known example. Since 1992, the magazine analyzes all articles from high-quality newspapers, counting how often the word "recession" appears. Based on this indicator, *The Economist* has been one of the first ones to announce the coming recession in early 2001 [9].

Such early indicators seem to exist also for life-critical matters. A recent analysis [14] indicates the consumption of cough sirup to be an early indicator for a possible anthrax terrorist attack. The early symptoms of an anthrax contagion are perceived in first instance as a common cold. These very first days are crucial for healing, though. Early indicators are vital. The authors assert that a certain increase in over-the-counter sales of cough sirup in an area points with a high probability to an anthrax epidemic. Nowadays, sales are reported by drugstores on a weekly or monthly basis, which wouldn't be much help. If government agencies could have real-time data about such sales, an epidemic could be discovered during the decisive first days. A step further, if data about home medication consumption would be available, this would be a more direct indicator; for instance at night, when drugstores are closed.

3.5 Cruising in the backseat – Economies on autopilot

Despite all these economic advantages, there are many drawbacks that have to be taken into account when designing ubiquitous computing systems. We will not address here the already mentioned social concerns, especially those specific to privacy and technology. In this section, we will rather concentrate on potential economic drawbacks.

A first thing to note is that activating an automated control procedure such as an airplane autopilot, or car cruise control will improve system stability under 'normal' situations - after all, machines are much better than humans in dedicating all their attention to a otherwise boring task. However, there will always be situations not foreseen in the software that can have disastrous consequences if not handled by humans. The 1987 stock-market crash for example has been partly caused by newly deployed trading software [38] that was programmed to trigger a selling action when a certain pattern appeared in the daily stock fluctuations. Since all traders had similar software, a small variation in the trading patterns got greatly amplified as all copies of the software continued taking the market down by executing their predefined selling pattern.

In a ubiquitous computing economic environment, additional concerns arise. That is because in order to be most efficient, an instant-time economy must be trimmed to be as slim as possible. However, when under such circumstances a minor unforseen event happens, the consequences could quite easily be disastrous. In the case of supply-chain management, for example, the elimination of the bull-whip-effect through transparent information flow along the supply-chain allows stocks to be dramatically reduced. But since all links in the chain can now reduce unnecessary stock to a minimum, any malfunctioning of the weakest link in the chain could potentially stop the supply across all partners.

The implications of ubiquitous computing deployment in business processes and economic transactions are thus both far reaching and diverse. The considerable potential savings of streamlined supply-chains, markets and private property will most certainly drive a large part of ubiquitous computing deployment in future years. Yet the greater the potential savings, it seems, the larger also the risk associated with sudden failures of such often complex and sensitive interaction models described above. The next section will try to describe some of the requirements for building *reliable* ubiquitous computing systems, and again look at the corresponding implications for system design that follow from them.

4 It depends – Building reliable ubiquitous computing systems

Today computers have already become irreplaceable aides for all kinds of settings and situations: embedded processors assist to perform critical medical operations, monitor patient's health conditions, automatically regulate temperature/ventilation in buildings or tunnels, and safely guide planes during landing or take-off.

Yet as we move closer to realize Marc Weiser's vision of ubiquitous computer [46], an ever increasing number of electronic devices, processors and micro-controllers are embedded in everyday objects. As a consequence, we may become more and more dependent on the proper operation of such technologically enhanced artifacts and environments, even if we are not fully aware of it. What is more, in a thoroughly computerized future, we might not even have the means to avoid such an ever increasing dependency.

The following sections will re-examine the implications of such a thoroughly interconnected world by describing existing dependencies in our daily life and stipulating the additional requirements on the reliability of ubiquitous computing systems in the future.

4.1 Been there, done that - Dependencies in society

Ubiquitous computing imposes a new, but certainly not the first dependency on our society. Already in primeval times mankind had to cope with dependencies that had their origin in natural environmental conditions or bodily restrictions: life depended for example strongly on natural forces and phenomena, such as the sun, the wind, and the rain, and the range of activity was typically limited to walking distances.

As human beings started to develop tools and shape their environment to their needs, they managed gradually to rid themselves of certain dependencies of old while at the same time introducing new, artificial ones. With the development of irrigation techniques, for instance, fields would yield higher crops and thus support a population greater in number, reducing the dependency on unpredictable weather conditions. Yet at the same time the well-being of the increased population would depend on the proper functioning of the irrigation methods in return. *Dependability* has consequently been an important issue already back then: If an irrigation technique would not meet *user expectations* in the long run, it would it would have been of little use and surely be abandoned soon.

Even though the concept of dependability as the "trustworthiness of a computer system such that reliance can justifiable be placed on the service it delivers" [22] was designed with distributed computer systems in mind, it apparently is an universal concept that also holds true for other technical systems such as irrigation methods.

Ubiquitous computing systems, too, will need to posses such dependability. And with respect to the vision of invisible computing, *meeting user expectations* is certainly another important aspect [36]. However, the implications we have seen in the previous sections also introduce additional requirements, such as persistence, manageability, control, and accountability. We will examine each of these requirements in the following sections.

4.2 Trust me – Relying on a ubiquitous computing environment

The vision of ubiquitous computing describes a system that resides in the background and unobtrusively provides catering to our needs. Since our needs change over time, depending on a large variety of circumstances, the system will necessarily need to dynamically adapt to various situations.

An example of such dynamic change has been given in the previous chapter, where dynamic pricing was able to find the best price depending on the availability of goods and their current demand. In order to have humans participate in such a dynamic market, however, both *persistance*, i.e., a certain inertia of the system that allows humans to react, and *manageability*, i.e., the ability to configure such dynamic pricing schemes to suit for example a certain taxation plan.

In order to lower the demands on human intervention in such a dynamic world then is the concept of *delegation of control*, where we put automated processes in control of otherwise boring routines, yet provide *accountability mechanisms* that allow us to understand complicated control flows.

When taken together, these requirements give raise to number of questions when following our pattern of finding real-world implications, which will be examined in the following paragraphs.

- **Dependability**: In ubiquitous computing, more and more devices become ever smaller in size and limited in the amount of available resources. And with an ever growing number of devices and appliances, the probability of failure for any single device increases proportionally, too. The limited amount of available resources leads to reduced device capabilities and more stringent resource restrictions. Especially energy becomes a critical resource of small self-contained devices. Besides, when space and resources are limited and energy consumption has to be kept low, there's little room left for hardware redundancy within a single device. Also, a user is likely to only possess a singe unit of a certain device type, e.g. just one personal digital assistant (PDA), one smart wristwatch, one digital camera or one key ring. Still it is highly desirable to achieve a high degree of robustness and fault tolerance. In this case, if there's only a low degree of hardware replication in the system, e.g.

a low number of duplicate devices, the threat of service disruptions due to device failures may be overcome by supporting *diversification* of system functions on different levels of abstraction. Diversification in the sense that the system is designed to be in a position to perform an operation or task in different, independent ways that have a minimum of hardware and software resources in common. E.g., to diversify a communications link the system should provide independent means of communication such as GSM [15], Wireless LAN, and infrared communications capabilities. Moreover, to diversify the accessibility of a service or of information, various independent gateways should be supported, e.g. a WAP portal [44] for mobile phones, a client application for PDAs or Laptops and a Web interface for manual access using a Web browser.

- **Predictability**: We are surrounded by a multitude of technical devices and infrastructures such as the telephone or electricity already today, yet their use is rather straightforward and *predictable*. Generally speaking: if you pick up the phone, you expect to hear the dial tone, and if this is not the case, you know that something went wrong, that there is a problem. However, this no longer holds for typical ubiquitous computing systems in general. Here, the ideal of the invisible, altogether unostentatious computer that silently hides in the background, might complicate or even impede the predicability of the system. If I am not aware of what's going on, I might not be in a position to notice the presence of failures, either. Knowing *what went wrong* and *when* may be essential, otherwise I might have no means to respond properly. E.g., a risk patient who cannot detect the failure of his continuous patient monitoring device will be deprived of live saving medical treatment in case of an emergency. Additionally, *excessive customization* has a similiar effect in reducing the predictability of a ubiquitous computing system.
- **Persistence**: In a world governed by the principle of total dynamics, what is valid in one moment may not hold any more in the next. And if the state of the world is no longer *persistent* even for short periods of time, if the rate of change in everyday matters suddenly surpasses our ability to adapt, human beings might face serious problems to cope with an excessive rate of changes. Such a phenomenon could also be described as a *lack of inertia* of the system in the sense that certain conditions that used to change only rarely in the past all of a sudden lose their constancy and begin to vary frequently. If information become obsolete the moment I have picked them up, it might become very difficult to gain valuable and valid experiences and know-how about the world around me to base my judgement and reasoning on; our human capabilities to readapt might simply be overstrained. In the long run this could have serious effects, e.g. cause a degradation of lasting experiences and thus add to a higher degree of uncertainty and disorientation in our human society. E.g., typically the costs of goods in supermarkets or the fares in public transport change rather seldom – by-and-by, we may memorize them fairly well if we wish to. In contrast, in the case of *dynamic pricing* it might become nearly impossible to foretell a future price.
- **Comprehensibility** and **Manageability**: As myriads of things become smart by means of embedded processors, memory and communication skills and are thus equipped with a life of their own, scalability becomes an issue. How does a world of tiny interoperating objects scale, and more, how can it be kept in check? E.g.,

assume that, in a supermarket, everyday life articles begin to dynamically set their own price according to demand, age and supplies available on stock. It is questionable to which extent these goods would still be manageable. It may become hard if not impossible to asses the value of all stock at one point of time, obstructing a reliable stock-taking process.

- **Control**: One major goal of smart environments and smart objects is to provide new means to unobtrusively assist us and hide the complexity of a technology-permeated world in order to improve the quality of our everyday life. Yet there is a fine line between smartness and vexatiousness, between just being helpful and anticipating or being headstrong and commanding. When should a smart device obey human orders and when follow its own line of action? While driving a modern car with anti-skid system on an ice-frozen street, for instance, you might gladly accept the smart brakes to intervene and prevent the wheels from locking, thus averting the danger of skidding. However, imagine your smart car detects that you stopped illegally within a no parking zone and therefore denies to open the doors. As a consequence, you may be pretty annoyed, but you'll probably accept that the car has, as a matter of principle, come to the right decision. Now consider a slightly different situation, that you are in a state of emergency and you rightfully stop by at the hospital entrance in the prohibited area, but the car hinders you from stepping out. The idea to have the car behave smartly and decide when to open the doors may be based on good intensions. But in a state of emergency this concept might go awry, leading to an involuntary *incapacitation of the user*; the user should have been in charge but lacked *control*. This may serve as an example that there might always be certain situations where full remote control or full automation may be harmful and counter-productive, calling for some means of manual override and user control.
- **Accountability**: One may think of business models that allow the short-time leasing of everyday life articles. Imagine leasing each seat one chooses to sit on during the day rather than paying entrance fees during a concert or buying a ticket for the train instead. During the day you'd automatically conclude leasing agreements to spend micropayments according to the price category and occupation time of a seat, no matter if you take a seat in public transport, in a cafe or even at your home. Looking at such a pricing scheme, verifying the money spent at the end of the month would become all but straightforward. Tracing back and verifying hundreds of transactions, micropayments or microleases is not only cumbersome and unrealistic to perform, but it raises also the question of accountability. E.g., there need to be mechanisms to a) decide and prove who has to pay money to whom or b) allow me to fend off illegitimate claims.

The above requirements represent another set of implications for ubiquitous computing. They are probably less obvious than popular privacy concerns or monetary economic aspects, but not less central, as they fundamentally influence the *acceptance* of ubiquitous computing technology. A ubiquitous computing system lacking those features will in most cases suffer from limited confidence of its users. Consequently, the deployed architectures will either not be used by the main target group (e.g., by the general public), or they have to be installed against popular will.

The next section will describe examples of such acceptance problems by looking at early ubiquitous computing critique, as it provides yet another view onto set of possible real-world implications that will be beneficial for a design that takes the concerns of the average citizen into account.

5 Critique of ubiquitous computing

It is not surprising – given the wide variety of far-reaching implications described in the previous three sections – that ubiquitous computing receives ample criticism from both scientists and the average citizen. Analyzing such criticism not only allows for better understanding current fears and misconceptions about ubiquitous computing, but may also add yet another view into the real-world implications of ubiquitous computing.

Ubiquitous computing criticism in today's literature can roughly be categorized into three different classes:

- criticism focusing on the *vision and goal* of ubiquitous computing, especially its all-encompassing nature and life-time coverage,
- criticism of the (negative) *effects* of ubiquitous computing, for example its already outlined surveillance and privacy threats, and
- criticism questioning the *cost-benefit* ratio of ubiquitous computing, stating that even if positive effects were possible, it might still not be worth deploying.

The following sections will examine each aspect in detail and try to summarize the existing literature, attempting to identify some of the implications that were found in the previous sections.

5.1 A better place? – Goals and visions

Mark Weiser describes ubiquitous computing as a technology that opens up enhanced ways of interacting with the world, based on the principle that computers themselves vanish into the background and are woven "into the fabric of everyday life until they are indistinguishable from it." [46]. Critics in contrast, often see ubiquitous computing as "an attempt at a violent technological penetration of life" [3], as "the feverish dream of spooks and spies – to plant a 'bug' every object – enlarged and re-shaped" [41], or even as a "project that aims at totality and that also stands close to the totalitarian" [1]. Why does ubiquitous computing meet such strong objections?

Far-reaching goals. One first and obvious answer to this question is grounded in the vision itself. The explicit goal of ubiquitous computing is to have a significant impact on *all* aspects of the existence of *every* human being in our society, by means of a revolutionary transformation of everyday life. It aims at the penetration of our the everyday environment with ubiquitous computing technology, in which "each person is continually interacting with hundreds of nearby wirelessly interconnected computers." [47]. In other areas of computer science a transformation of society was not an explicit goal, but rather a side effect that happend by chance and was not planned beforehand. Ubiquitous computing in contrast proposes an all-embracing computerization of society in its inception, one that aims at the totality of things and beings, letting it appear as a "project that aims at totality" [1]. In the vision as formulated by Marc Weiser, every thing, place, and being is affected, leaving no freedom of choice to join in or opt out.

Vagueness of the ubiquitous computing vision. Many critics [1, 3, 25] argue that ubiquitous computing scenarios appear to be pretty vague compared to the enormous research efforts required and the dimensions of the envisioned goal. The world will "somehow" be made "smart", but the "somehow" is not specified further. Technological advances – like miniaturization, increasing processing power, and wireless connectivity – open up new application possibilities, but it is as yet unclear how to use these possibilities constructively. "Everything will be connected to everything else", but "no one has any idea what all those connections will mean" [25]. There is a gap between technological possibilities and our ability to put them to good use, or even to assess their benefit. This divergence between technological intensification and decreasing perceived value is what John Thackara [42] calls the *innovation dilemma*: "We know *how* to make amazing things, but we don't know *what* to make", or more pointed: "We are brilliant on means, but pretty hopeless on ends". Maybe it is important here, to formulate a more concrete vision of ubiquitous computing, one that features worthy, valuable, realistic goals in order to gain widespread acceptance.

5.2 But what is it good for? – Effects

Even though many of the possible effects of ubiquitous computing are not yet clear, it is relatively easy to imagine general negative consequences that such developments might have. Its core categories fall along the liens identified in the previous three sections.

Danger of surveillance and loss of privacy. Not surprisingly, the most direct fear associated with ubiquitous computing is that its mechanisms could be misused for efficient and merciless universal surveillance and lead to a degradation or complete loss of privacy. To quote R. Lucky [25]: "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." A more differentiated view on this topic has been presented in section 2.

False promises. Winner [51] argues that ubiquitous computing raises false expectations if it promises that its realization will simplify our lives, save time, and liberate ourselves from toil. For him this has been a recurring claim of consumer technology throughout the twentieth century. He cites anthropological studies on Silicon Valley employees that show that these people have "an endlessly busy, complicated, precariously balanced, strung-out existence in which traditional boundaries between work and leisure have evaporated" [51]. He says that in such a situation "adding smart machines to every corner of the built environment does nothing to alleviate these patterns of hurry, stress, and disconnection from people" [51]. The application of ubiquitous computing technologies would not lead to more spare time or more relaxed lives, but instead enables people to be more effective in the activities they perform. Looking at the economic implications outlined in section 3, it is easy to imagine that many of the driving factors behind future ubiquitous computing systems might very well by motivated by increased productivity and profits rather than simplifying our lives.

Loss of control. An important aspect of the design of ubiquitous computing systems is the feeling of being in control over one's surroundings and the "loyalty" of things in them. If, for example, the refrigerator refuses to open after having talked to the bathroom scale, or if my car refuses to turn off or to let me open the door, because I want to park illegally, then one easily feels "surrounded by enemies and traitors" [25]. The car and the bathroom scale, being part of a ubiquitous computing network, might not be completely "loyal" to their owners, but to the insurance company, the law, or the manufacturer. As alluded to already in sections 3 and 4, true ownership of goods might indeed be replaced by licensing models that only allow the use of certain capabilities without providing complete control over them – a concept that might further increase suspicion towards a future that people can rely on.

5.3 Net worth – Perceived value vs. social costs

Beyond ubiquitous computing visions and its actual effects, the *cost-benefit* analysis of such changes, even if positive in nature, is also often questioned. This is because minor beneficial gains through the use of ubiquitous computing might be set off for example by huge monetary spending required for research, or through more important negative side-effects.

Marginal perceived value. According to Araya and Winner, the needs presented in many ubiquitous computing scenarios have a marginal character. They do not enable things that were impossible before, but only gradually enhance human activities. Weiser [46] states that "ubiquitous computing will enable nothing fundamentally new, but by making everything faster and easier to do, with less strain and mental gymnastics, it will transform what is apparently possible". From Araya's [3] point of view, this does not justify the enormous research efforts and the complexity of the required infrastructure. For him, ubiquitous computing has little to do with significant human needs, but with the unfolding of technology per se. He calls this apparent primacy of technology over human needs *technological absolutism*. The "average citizen" might like to agree: "Descriptions of the world of ubiquitous computing are dazzling, if only for their sheer silliness. If you rate humanity's needs for the coming century on a scale of 1 to 10, none of the products and services depicted in Levy's article rises much beyond a score of 1.5," writes one reader of an article on ubiquitous computing in Newsweek [30]

Change in human–world relationship. Philosophical analyses of ubiquitous computing consider its effects on the human–world relationship. Ubiquitous computing fundamentally changes the environment in which we live. It becomes "intimately tuned to us" [3]. The physical surroundings become an extension of our bodies, while ubiquitous computing extends our nervous system, with sensors acting as artificial nerves. The environment becomes a "subservient artifact" [3], it becomes *us* rather than *other*. Adamowsky [1] considers the question if we are really willing to live without an *outside* as fundamental. According to Adamowsky, the *outside* will disappear and we will live in *set* worlds. These *settings* concern aspects of the real world that are mapped onto the digital world and are realized as models, simulations, and virtual counterparts. Araya calls these settings *digital surrogates* and regards them as characteristic for ubiquitous computing: The inability to disseminate the physical world requires us to disseminate digital surrogates of the world and results in a transformation, displacement, substitution, and loss of fundamental properties of the world.

5.4 Do the right thing – Implications of criticism

The realization of ubiquitous computing has far reaching effects for each individual and society as a whole. Therefore, “the ubiquitous computing proposals should not remain unchallenged, but be subject to intense investigation” [3]. Can the vision of ubiquitous computing for the future development of society be based solely on technological means? Does ubiquitous computing – if ever put in place – really make the world a better place? Does it address the right issues, i.e. issues that can be solved by technological means? Do we really want to spread sensors, computation and connectivity to everything man-made and even nature-made?

Ubiquitous computing appears to be driven by technology and put into practice by “technologists” who have neither the legitimacy to design and transform the everyday environment of human beings, nor have proven their competence in doing so. What the critique of ubiquitous computing presented here seems to make clear is the need for a public debate about the goals and ideas set forth in ubiquitous computing. As long as no such debate is taking place, ubiquitous computing will easily be misinterpreted, resulting in irrational fears and resentments. An example is that the guiding priciple of ubiquitous computing – to make computers vanish into the background – is by some critics seen as an attempt to let ubiquitous computing go unnoticed in order to bypass resistance against it. This view is sustained by a quote by Marc Weiser, if taken in isolation: “the most profound revolutions are not the ones trumpeted by pundits, but those that sneak in when we are not looking.” [48]

6 A brave new world?

The deployment of ubiquitous computing systems in the real-world will in many cases have implications beyond the technically obvious ones. Whether it be personal privacy, national economies, or social acceptance – designers of ubiquitous computing systems can greatly benefit from evaluating the effects of putting ubiquitous computing into the real-world using existing concepts in disciplines such as social, economic, and legal sciences.

As difficult and in vain predictions of the future often are, our previous discussions nevertheless allows for the identification of a variety of implications, should large-scale ubiquitous computing installations be taken into the real-world: values and motivations change; personal border are crossed due to monitoring and searching; business methods provide increased profits at the expense of safety margins; economies accelerate and rewrite social values; confidence in our environment might suffer; and our attitude to the world that surrounds us might significantly change.

“Science Finds - Industry Applies - Man Conforms.” This motto of the Chicago World Fair of 1933 maybe exemplifies the long way scientists and engineers have come since then. Ubiquitous computing will require us to negate such kind of thinking even more, by explicitly focusing on human and societal needs before applying any of the concepts computer science has found. The growing field of *engineering ethics* [28] might serve as a guidance for designing particular applications of ubiquitous computing. Maybe by applying the ethics of utility theory and Kant’s categorical imperative, together with a proper analysis along social and economic theories, can we hope to find robust design methodologies for building and deploying *acceptable* ubiquitous computing systems in the future.

References

1. Natascha Adamowsky. Kulturelle Relevanz. Ladenburger Diskurs “Ubiquitous Computing”. Available at: www.inf.ethz.ch/vs/events/slides/adamowldbg.pdf, February 2000.
2. Philip E. Agre and Marc Rotenberg, editors. *Technology and Privacy: The New Landscape*. The MIT Press, 1998.
3. Agustin A. Araya. Questioning ubiquitous computing. In *Proceedings of the 1995 ACM 23rd annual conference on Computer science*. ACM Press, 1995. Available at: doi.acm.org/10.1145/259526.259560.
4. Victoria Bellotti and A. Sellen. Design for privacy in ubiquitous computing environments. In *Proc. of the European Conference on Computer-Supported Cooperative Work*, 1993.
5. Bluetooth Consortium. The Bluetooth Project Homepage. www.bluetooth.com/, 1999.
6. David Brin. *The Transparent Society*. Perseus Books, Reading MA, 1998.
7. Vannevar Bush. As we may think. *Atlantic Monthly*, July 1945.
8. Peter Cochrane. Head to Head. *Sovereign Magazine*, pages 56–57, spring 2000.
9. Rrrrrrrrecession? *The Economist*, January 4, 2001.
10. Amitai Etzioni. *The Limits of Privacy*. Basic Books, New York NY, 1999.
11. Simson Garfinkel. *Database Nation*. O’Reilly, Sebastopol CA, 2000.
12. Robert Gellman. Does privacy law work? In Agre and Rotenberg [2], chapter 7, pages 193–218.
13. W. Wayt Gibbs. As we may live. *Scientific American*, November 2000.
14. Anna Goldenberg, Galit Shmueli, Richard A. Caruana, and Stephen E. Fienberg. Early statistical detection of anthrax outbreaks by tracking over-the-counter medication sales. *Proceedings of the National Academy of Sciences of the United States of America*, 99(8):5237–5240, April 2002.
15. GSM Association Homepage. www.gsmworld.com/index.shtml.
16. Hitachi, Ltd. The mu-chip. Available at: www.hitachi.com/products/material/rfid/.
17. Java Card Technology. java.sun.com/products/javacard/.
18. Robert O’Harrow Jr. Prozac maker reveals patient e-mail addresses. *The Washington Post*, July 4, 2001.
19. David M. Kreps. *A Course in Microeconomic Theory*. Princeton University Press, 1990.
20. Accenture Technology Labs. Silent commerce applications. Available at www.accenture.com/xd/xd.asp?it=enWeb\&xd=Services\%5CTechnology\%5Ctech%_payperuse.xml, 2001.
21. Marc Langheinrich. Privacy by design - principles of privacy-aware ubiquitous systems. In *Proceedings of Ubicomp 2001*, pages 273–291, September 2001.
22. J. C. Laprie. *Dependability: Basic concepts and terminology in English, French, German, Italian, and Japanese*. Dependable computing and fault-tolerant systems, v. 5. Springer-Verlag, Wien ; New York, 1992.
23. Hau L. Lee, V. Padmanabhan, and Seungjin Whang. The bullwhip effect in supply chains. *MIT Sloan Management Review*, 38(3):93–102, Spring 1997.
24. Lawrence Lessig. *Code and other Laws of Cyberspace*. Basic Books, New York NY, 1999.
25. Robert Lucky. Everything will be connected to everything else. *Connections. IEEE Spectrum*, March 1999. Available at: www.argreenhouse.com/papers/rlucky/spectrum/connect.shtml.

26. Gary T. Marx. Murky conceptual waters: The public and the private. *Ethics and Information Technology*, 3(3):157–169, 2001.
27. Robert N. Mayo. The factoids project. Available at www.research.compaq.com/wrl/techreports/abstracts/TN-60.html.
28. Gene Moriarty. Three kinds of ethics for three kinds of engineers. *IEEE Technology and Society Magazine*, 20(3):31–38, fall 2001.
29. Kristine S. Nagel. Family intercom: Developing a context-aware audio communication system. Presentation at Ubicomp 2001. Available at www.cc.gatech.edu/∼kris/research/intcomm/ubi1/sld001.htm, September 2001.
30. Newsweek. Griping about gadgets – letters to the editor, 1999.
31. Council of Europe. Convention on cybercrime. Available at: conventions.coe.int/Treaty/EN/projets/FinalCybercrime.htm, November 2001.
32. Joseph P. Pickett, editor. *The American Heritage College Dictionary*. Houghton Mifflin Co, 4th edition, April 2002.
33. Bradley Rhodes. The wearable remembrance agent: A system for augmented memory. *Personal Technologies Journal. Special Issue on Wearable Computing*, 1:218–224, 1997.
34. Bradley J. Rhodes, Nelson Minar, and Josh Weaver. Wearable computing meets ubiquitous computing - reaping the best of both worlds. In *Proc. of the third International Symposium on Wearable Computers (ISWC' 99)*, pages 141–149, San Francisco, CA, October 1999.
35. Mark Rotenberg. Testimony and statement for the record. hearing on privacy in the commercial world. Available at www.epic.org/privacy/testimony_0301.html, March 2001.
36. M. Satyanarayanan. Pervasive computing: Vision and challenges. *Ieee Personal Communications*, 8(4):10–17, 2001.
37. M. Satyanarayanan. A catalyst for mobile and ubiquitous computing. *IEEE Pervasive Computing Magazine*, pages 2–5, January 2002.
38. Ludwig Siegele. How about now? A survey of the real-time economy. *The Economist*, pages 3–18, February 2, 2002.
39. Polly Sprenger. Sun on privacy: 'get over it'. January 1999.
40. Stunz. Privacy's problem and the law of criminal procedure. As cited in [24].
41. Steve Talbott. The trouble with ubiquitous technology pushers, or: Why we'd be better off without the MIT Media Lab. *NETFUTURE: Technology and Human Responsibility.*, January 2000.
42. John Thackara. The design challenge of pervasive computing. *Interactions*, 8(3):46–52, May/June 2001.
43. United States of America. The Declaration of Independence and the Constitution of the United States, August 1998.
44. Wireless Application Protocol Forum. www.wapforum.org.
45. Samuel Warren and Louis Brandeis. The right to privacy. *Harvard Law Review*, 4:193 – 220, 1890.
46. Mark Weiser. The computer for the 21st century. *Scientific American*, pages 94–104, September 1991.
47. Mark Weiser. Some computer science problems in ubiquitous computing. *Communications of the ACM*, 36(7):75–84, July 1993.
48. Mark Weiser. Ubiquitous computing. *IEEE Computer*, 26(10):71–72, 1993.
49. Bo Westerlund, Sinna Lindquist, and Yngve Sundblad. Cooperative design of communication support for and with families in Stockholm. Available at interliving.kth.se/papers.html, September 2001.
50. Alan F. Westin. *Privacy and Freedom*. Atheneum, New York NY, 1967.
51. Langdon Winner. The voluntary complexity movement. *NETFUTURE: Technology and Human Responsibility*, September 1999.