Exploring the Business and Social Impacts of Pervasive Computing

March 2006

IBM Zurich Research Laboratory
Swiss Re Centre for Global Dialogue
TA - SWISS

Exploring the Business and Social Impacts of Pervasive Computing

March 2006

Swiss Re
Centre for Global Dialogue

IBM

TA SWISS

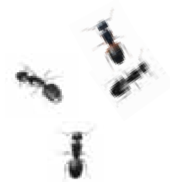

Foreword

Our world is becoming increasingly filled with small processors, embedded in everyday objects interacting with one other via the digital world of the Internet, following their rules and adapting to their surroundings, all without human intervention.

The development of *pervasive computing* has the potential to cause a paradigm shift in how societies apply and think of technology.

For this reason, decision makers need not only to understand and enable the use of this new generation of technologies, but also prevent its potential misuse. Privacy becomes central in the debate on these new technologies. Does *pervasive computing* require a step change in privacy policy? Does it also require a new generation of privacy technologies?

A team of interdisciplinary experts representing our three organisations has been working together during the last year to identify and explore the central issues and consequences of *pervasive computing* on the economy and society in general.

We hope the views and findings contained herein will provide impulses for further discussions with us and others.

Fritz Gutbrodt
Head of Centre for Global Dialogue, Swiss Re

Krishna Nathan
VP Services and Director, IBM Zurich Research Laboratories

Sergio Bellucci
Managing Director, TA Swiss

Table of Contents

Foreword
The Ant's Foreword
Introduction

Section 1
Describing *Pervasive Computing*
1.1 Connectivity
1.2 Embedded processors
1.3 Invisibility
1.4 Location and mediation
1.5 Clustering and the drive to order

Section 2
Essential Infrastructure
2.1 Introduction
2.2 Devices
2.3 User interfaces
2.4 Network infrastructure
2.5 Standards and enabling incrementalism
2.6 Network traffic

Section 3
Unpredictability, Collaboration and Optimisation
3.1 System fragility and criticality
3.2 Designing systems and unpredictability
3.3 Who or what is the circuit breaker?
3.4 Optimisation
3.5 What will the increase in data and analysis do to us?

Section 4
Underlying Social Trends Affected by *Pervasive Computing*

Section 5
Social Impact of *Pervasive Computing* and Privacy
5.1 Introduction
5.2 Shifting boundaries of privacy
5.3 *Pervasive computing* "end-game"
5.4 Will *pervasive computing* really make such a difference?
5.5 How to become a consenting individual?

The Ant's Foreword

At first glance it may be a surprise that an ant has been invited to write the foreword for a report produced by a group of academics, risk experts and industry researchers. Yet this is not the first time an ant has been used to describe some of the outcomes of complex, self-organising systems. Ants interact locally with each other according to some basic, simple rules and, at the same time, serve a highly organised colony capable of surviving many years longer than the lives of its inhabitants, all without centralised control.

This form of organisation makes the ant popular when trying to clarify the type of environment envisaged by pervasive computing. First of all there is the sheer number of processors and appliances that are going to be connected together in the near future: exponential growth rates suggest that the few billion devices we have today could turn into a trillion inside 20 years. Many of these machines will not be large and complicated as are, for example, the personal computers of today. Instead they will be small sensors and tags attached to objects, feeding reams of data to other larger machines within a network and receiving a small number of commands in return. So it is not so much that individual objects begin to think, but that each starts to play a small role within a broader system, triggering activity beyond their immediate sphere of influence and enabling new properties to emerge at the level of the system, just as ants manage to do at the level of the colony.

I am also here to remind the reader of how quickly humans jump to conclusions, interpreting the world as if they are at the centre of everything. When I follow a trail to gather food, it is immediately assumed I know what I am doing, that I have a purpose in mind, like a soldier sent into battle by their queen. The truth is I do not have a clue what is going on and if I did I would not be as effective. In the same way, when I say that with pervasive computing humans become part of the system and are always observable, via data systematically collected and clustered around their identity, then the assumption is that there will also be someone there to view it. And the idea that objects interacting with digital information will generate a new drive to order within the chaotic information sphere will lead, no doubt, to murmurings of self-aware computers, if not a new global consciousness. So I, Antigone, the lone ant shall act as a reminder of how quickly humans run up against the limits of reality, slipping into the seamless expanse of their imaginations, assigning purpose and intentional meaning wherever they go.

That said, it does seem as though a new computing era is beginning, a shift from man-machine interaction to machine-to-machine collaboration: tasks performed in the background, objects getting on with organising things, tiny sensors feeding software and information systems with multiple feedback loops. Like an Internet of things. Humans will no longer have to lift a finger, enjoying more time in their deck chairs, the sun shining and the drinks cooled, at least for some.

But of course, if something like pervasive computing is on the horizon, there will also be concerns. Never mind the list of benefits from the visionaries, what is this new technology? How and when is it going to get implemented? Of course if an ant wants to get a snapshot of what most concerns a group of humans, it will look at what they project into the future: scenario analysis techniques, mapping the future, envisaging how things will change and what applications will be used. This must be one of the big differences between ants and humans, this need to see the "big picture", some way of understanding what is going on, even if the capacity to influence it is quite slight.

From my perspective, however, this project seems to be different. The experts have left their crystal balls at home and it has been accepted that the next ten to twenty years are going to be a roller coaster ride leading to either a piecemeal advance in working with computers or a major phase transition in social and economic structures. The question is then, whether pervasive computing is part of a paradigm shift: a change in perspective that lets old problems be seen in a new light. Paradigm shifts do not generally happen overnight and, while often based in fundamental scientific research, they also directly influence social and intellectual models for interpreting events. Examples would be how Freud's notion of the unconscious has dominated popular notions of the mind, or how quantum mechanics has made the theoretician aware that the individual observer is always already part of the experiment. Such paradigm shifts play an important role in shaping an ability to perceive and rationalise risks and benefits and therefore maintain trust and, potentially, engagement in social, political and economic structures.

Of course as an ant the biggest paradigm change I undergo is having my colony destroyed and having to spend a lot of time relearning my rules of interaction with other survivors. But I recognise the pace of change humans undergo is increasing, especially in technology implementation and it gets more difficult to work out when a fundamental shift in perspective is occurring or whether there is simply a new gadget on the shelf.

So here it is. Let the ants bustle across the pages of this book. They are trying to build a thread of thought that is not so much about predicting the future as increasing awareness of the possible parameters in which changes might occur.

With the kindest regards

Antigone, the ant

Introduction

This text has emerged from a series of workshops held during 2005 to explore the potential social and economic impacts of *pervasive computing*. The participants were drawn from the international networks of the IBM Research Laboratories, the Swiss Re Centre for Global Dialogue, and TA-SWISS. They collaborated to create both this text and the agenda for the conference on *Pervasive Computing* held at the Swiss Re Centre for Global Dialogue in September 2005. The workshops focused on finding an approach to *pervasive computing*, defining its key characteristics and setting a context within which to explore its social, economic and personal ramifications.

A single author, Lawrence Kenny, was charged with developing the material from each consecutive workshop and establishing a platform for an ongoing conversation, the results of which are captured here.

While often appreciated and helpful to participants, a cross-discipline dialogue poses the challenge of how to document the flow of ideas that emerge. Such dialogues often take topics which are not yet well-defined, bringing together conflicting views, generating new insights and innovation, prior to the rigorous research that emerges from a single expert community. Such dialogues should indeed be a risky process, with participants reaching out beyond their own areas of expertise and any resulting text must endeavour to reflect this without sacrificing the demand to communicate an understanding of the core issue.

This text is a response to this challenge, which seeks to create an approachable overview out of the soup of factors that influence *pervasive computing*, whilst avoiding an overly simplistic and over-arching definition of the subject, which at this early stage of its development could be highly misleading.

The first section describes some of the characteristics thought to be important to *pervasive computing*, cumulatively defining the key aspects of the subject. The second section is a more concrete explanation of the areas of essential infrastructure required to enable the technology. The third section explores ways in which this new technology extends areas of potential fragility in how we design and use complex systems. The fourth section begins to describe more directly the social impacts of these enabling and systemic issues. Each element described may appear modest in itself, but, when combined, areas do emerge where *pervasive computing* could present major social change.

Finally in the last section the issue of privacy is explored in more detail. Workshop discussions repeatedly returned to this issue and it has therefore been singled out for further consideration. There is a deliberate avoidance of repeating the histories of early privacy conflicts, for example, how major companies plan to use RFID to gain further insight into the behaviour of their customers. Instead the focus rests on why it is that privacy issues act as a pressure point for social unease in technology developments and what types of changes may occur to the accepted framework of ideas on privacy issues.

Section 1: Describing Pervasive Computing

Forget the icon of the white, box-shaped computer, think networked processors, collaborating with each other to perform defined tasks, without human intervention. This is *pervasive computing*. It sees the next stage of man-machine interactivity as letting the machines get on with things themselves, collaborating with each other and being able to understand the different contexts in which they operate. It is a challenging vision that moves beyond current models of human interaction with computers and communication devices. This chapter seeks to explore some of the key characteristics that lie behind this vision.

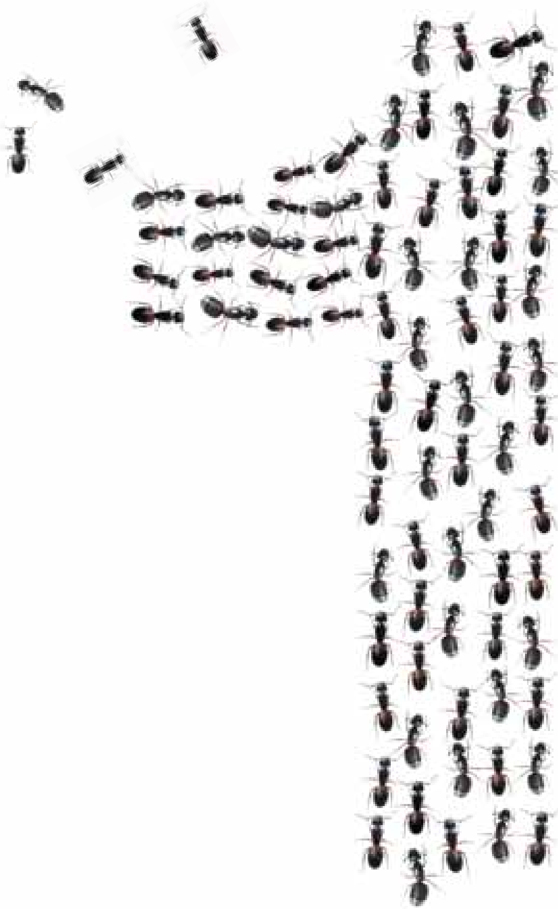

1.1 Connectivity

Connecting things generates the potential for new behaviours and applications. That is what happened when villages, cities and countries were connected via transport networks and it also applies to communication devices for mobile telephony and the Internet. And, like the ant colony, connecting things together into self-organising systems generates the potential for new properties to emerge at the system level, which cannot be predicted from the constituent parts.

Pervasive computing is then, a further extension of this process of plugging things together, in this case the millions of computer processing units residing in households, transport, industry and commerce, and the new generation of sensors and electronic tags that are slowly filtering into our social environment. These devices are often relatively straightforward processing units that collect pre-defined data on their environment and pass it to a controlling system that possibly sends a small number of basic commands in return. And, of course, once the data is fed from the devices to their controlling system, these can also be connected with other services and systems, creating yet another network. It is like a hall of mirrors reflecting an infinite regression towards the smallest of the small and the largest of the large.

An important factor in connecting these tiny sensors, tags and processors is their sheer number. Predictions suggest that in twenty years time there will be literally a trillion devices with the capacity to send and receive signals. This interactivity is not a virtual Babylon where all objects are attempting to communicate with each other all the time. For example, sensors buried in the concrete supports of a bridge, could remain dormant until disturbed by a tiny change in energy levels within the sub-atomic construction of the concrete, indicating the beginnings of a crack. It is rather that the potential is always there for objects to be connected so that interactivity becomes a ubiquitous property of our environment, something we take for granted.

Exactly which applications and new cultural behaviours are going to emerge from *pervasive computing* is difficult to predict. As an example, take the Short Message Service (SMS) used on the mobile phone. Originally planned for use by telephony engineers, it has become an important source of revenue and is used widely in many cultural environments, from organising teenage parties to remote adjustment of heating systems. For SMS, the unobtrusiveness and ease of use have proved more important than the inherent technological limitations, much to the surprise of the engineers.
In the absence of a “killer application” then, determining how technology will be taken up is problematic. Certainly, so-called intelligent fridges and other devices that simply automate everyday processes are unlikely to unlock the

investment required. More promising are new uses for health monitoring, security and improvements in supply chains. It is these specific types of environment or context that will first create *pervasive computing* as a phenomenon. For example:

1. Different computing devices will "talk" to each other, updating their position, switching networks, downloading relevant data
2. Objects defined through their role in consumer, military or industrial supply chains
3. Specific limited "smart" environments such as a workshop room or car (so-called mediated spaces)
4. Specific systemic environments such as a hospital, house or transport system
5. Human beings acting as "hosts" for wearable computing and health implants
6. Everyday objects in our environment such as pens, glasses, tables and toys

1.2 Embedded processors

Embedded processors are purpose-built, single function processors, capable of monitoring and reacting almost instantly to changes in their environment and taking corrective or responsive action. Over 90% of computer processing units operate in this way today, often requiring no attention from the user of the machine. They have long replaced mechanical or electrical devices because it is cheaper and simpler to use an embedded computer to run, for example, an alarm clock, than to actually build it using traditional engineering methods. Our reliance on such embedded processors means no transport system, household appliance or other device is complete without them. They perform their often highly repetitive tasks with the simplest possible algorithms and little reliance on any external services.

An important element here is that often the user does not directly control what an embedded system does. Indeed, that is their beauty, they operate as tools, tirelessly and without needing any human attention. In the case of safety systems this element of taking the control out of human hands is very explicit. Machines are simply quicker and more reliable than a human at performing repetitive monitoring and calculation tasks.

Given the role of embedded processors, it is clear computational devices already play an unseen and important role in our environment. With further advances in materials science, these devices will be made literally invisible to the human eye. It is these embedded processors which are developing secondary functions through their connection to networks, becoming an important basis for *pervasive computing*. As it is possible to not just embed processing units, but also to network them, information can be delivered on identity, location and other variables to other machines within a network. The information or software layer connects with classes of physical objects that can feed data on themselves and their environment and be acted upon in a pre-defined manner via embedded processors in close to real time.

Of course, the simplicity of this description hides much complexity. In particular it prioritises an aspect of processors, namely that they are connected to real objects that play a role within our daily life and environment. If *pervasive computing* presents a shift to processors that collaborate in the background without human involvement, will they always do so to fulfil human ends? This does not presume that the processors' independent activity is intelligent or that they can determine their own intentional actions, rather it asks whether we understand the full scope of the processors' world when we explain all its behaviours in terms of a tool to fulfil human needs?

Characteristics of *pervasive computing*

Small: Ongoing miniaturisation of components, moving to invisibility

Embedded: Components are placed on or within other devices, objects or living beings

Networked: Flexible capacity to exchange data and software components with other devices and platforms

Context sensitive: Collect and exchange data on their environment and the host object via sensors

Adaptive: Implement changes and modifications at the software and object level

Collaborative: Ability to discover other objects and interact with them to establish cooperation on the software or information level

Network Volume: Sufficient in number and regularity of interaction to create network behaviours

This list of characteristics is interesting in how unremarkable it is. It indicates the extent to which the trajectory of the technology used in *pervasive computing* has already been determined, even if the implementation remains problematic. In this sense, *pervasive computing* is not about the introduction of a single technology, but about a potential qualitative change that may arise through an increasingly integrated technological environment.

1.3 Invisibility

Is the computing device going to literally disappear and what does this mean? This question focuses on how the real and virtual environments are connected through the use of devices. This happens in several ways:

- Connecting the real world and the information world (e.g. sensor devices)
- Giving us access to the information world (e.g. personal computers, mobile phones etc.)
- Providing services between the information and real worlds (for example, credit cards and web services)

If devices are critical in making this connecting process happen in what sense are they going to disappear? The most obvious way in which this would occur is that the ongoing miniaturisation of processors allows them to become invisible to the naked eye and that these processors are embedded in other objects and thus hidden from view, as with modern braking systems in a car. There is, however, a second sense of invisibility, namely, that access to the computing environment becomes such an everyday part of life that there is no longer a noticeable distinction between the computing device and any other object with which humans interact.

The roots of *pervasive computing* are often traced to the early 1990s and, in particular, to the research programmes devised by Mark Weiser, Head of The Computer Science Lab at Xerox PARC. At that time there was a dissatisfaction within parts of the research community that computers demanded too much of the user's attention. Rather than being a means to an end, they tended to become the end itself, isolating the user and demanding they focus on the processing and procedures of the machine.

In 1990, the shift from mainframe computing to that of the personal computer or PC was already under way and the potential impact of hand-held mobile devices and the Internet was moving from vision to reality. Researchers therefore started to envisage how computers could recede into the background, responding to voice and gestural commands and with screens built into walls, desks or floors, encouraging collaboration between members of a team and not with the computer as such. The human would be at the centre of this environment, interacting with machines in natural language. The tyranny of making endless steering decisions with the click of a mouse would be at an end.

While many have found this goal, of treating technology as a means and not an end, compelling, stories of the death of the personal computer seem to be exaggerated. Nevertheless there remains a powerful analogy in treating computing devices like instruments.

The surgeon uses a scalpel to cut, but his attention is focused on the location of the incision and possibly on envisaging his next moves rather than on needing to be aware of the scalpel itself.

1.4 Location and mediation

As computation is embodied in objects and not just computers, location and the spatial environment will increase in importance. The ability to locate objects becomes a key enabling factor. Once objects can provide a self-description and report on specific attributes of their surroundings, it is possible to create a grid-like digital representation of the spatial and temporal environment and cluster further data around the identified object. This is different from the Internet age that championed a virtual reality in which space, beyond that generated on-screen, was irrelevant. Distance was eradicated as the physical location of the computer once it was logged onto a network was considered trivial. The world became "flat", without relief, fitting with the on-going globalisation of economies and cultures.

Pervasive computing comes about when multiple objects get connected with one another as part of a network. This does not mean that each independent embedded processor becomes super powerful by virtue of a wireless connection. Rather the objects become part of a network which means our experience of them is potentially even more mediated. There are physical intermediaries in the form of computer servers and other technical hardware, and there is an information intermediary as the data collected from a network is transposed with the addition of metadata and the application of software programming. And it can be argued there is a semantic and heuristic intermediary, as our needs are interpreted by the protocols and rules of engagement governing the information sphere.

This mediation of our environment, where there is global access to localised information, is an enormous opportunity for new services. The oft quoted example of a car giving background information as it passes through a city, indicating restaurants or a cinema according to pre-defined criteria, will be closer to fruition.

1.5 Clustering and the drive to order

A safe prediction is that *pervasive computing* will generate more data, a lot more. The volume grows each year, even before the deployment of more sensor devices. Optimistic researchers suggest that the storage and transmission of this data will not be problematic. But they are less sanguine about the procedures to aggregate it, parcelling it up into chunks that can be analysed and mined to generate meaningful information. This is an enormous task.

The sphere of digital information continues to expand and currently defies any universal ordering system. The number of Internet hosts is rapidly growing, in 1999 there were less than 100,000,000 and today there are over 400,000,000. The so-called “spiders” of search engine companies that go out and classify the information on these databases cannot keep up with this rate of expansion. So while the order and organisation in some areas may continue to improve, overall our information sphere is becoming more chaotic.

Will a change in the quantity of data collected create a qualitative change in what it can communicate? The simple answer is yes, if the data elements are distinct and correlations can be made between them. If the data elements share something in common, then there is a hook which facilitates the ordering and clustering of the data. One hypothesis is that the data collected using sensors is special because it deals with the physical world of space and time. Objects can identify themselves, can be located and “events” can be recorded, i. e. changes to specific attributes, through time. In a technical sense, these data elements are self-describing, which means they are a powerful hook on which to form data clusters and detect further correlations. This means the rather chaotic virtual world of information becomes inhabited by familiar objects and events that already carry meaning. Connecting objects in this sense generates a drive to order among the chaotic sphere of data.

What an object does is define a fixation point around which the data can cluster. The fixation point may also be an event, for example paying for an item in a department store: the credit card company receives information on the transaction, the maker of the object receives information that it has been purchased, influencing the production schedule and the department store chain receives information with regard to its stocking position. Each of these events has led to the creation of “collateral data”. It is of course not the information itself that is interesting, it is that there are millions of such activities going on concurrently which become data sets and combine for review and analysis. These data sets may well cluster in a self-organising manner, getting recruited to a cluster through interaction with another activity that shares a common element. This set may then combine with other clusters should it continue to recruit other data or it may simply dissolve back into the nebulous cloud of digital information.

Alice's party

Alice is an accountant, living in the garden flat of a newly built residence on the outskirts of town. She has decided to throw a housewarming party for friends and neighbours. Being shy and uncertain as to how many people would want to come, she has told her friends to also invite people should they wish to. The evening of the party has come and it has started with a bang. Her good friends arrived early and she has enjoyed taking coats, pouring drinks and generally making people feel at home. As the evening progresses, however, something changes. The party is crowded with people she does not know, friends of friends, helping themselves to drinks and changing the music. Someone is even washing up glasses. Alice finds herself first in the kitchen and then slipping out, alone, onto her newly laid patio, staring back through the picture window. A shiver of uncertainty passes through her as she wonders how she had ended up outside, looking in at her own party.

What should Alice do? Call the fire brigade on her mobile and shout "fire" to break up the party? Or trust that everything will be alright, treating the whole event as a tremendous, swinging success? The former means she gets control of her flat back but offends the neighbours and loses some potential friends. The latter, that she gives the role of policing the party to the group, in the hope that their rather effective self-organisation will dissolve at the end of the party as quickly as it has emerged.

Alice's concerns may continue to reverberate through this discussion on pervasive computing. How should one respond as part of a network where simple activities can lead to the emergence of rather intelligent behaviour and order at the system level? And how does this influence the individual's feelings of control and expectations of the predictability of the environment around us?

Section 2: Essential Infrastructure

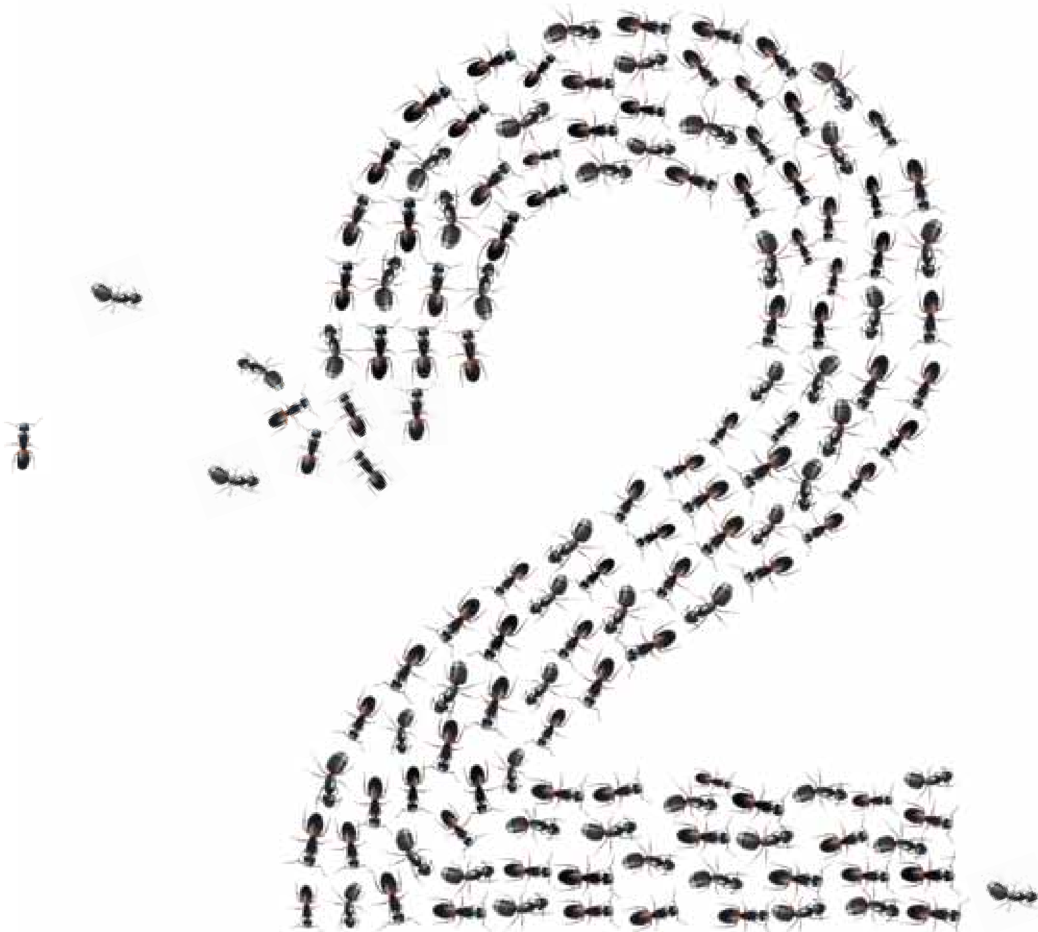

2.1 Introduction

Without the appropriate hardware devices and networking infrastructure, the *pervasive computing* vision will fail. It needs a computing infrastructure that seamlessly helps people in accomplishing their tasks, whilst making most of the actual computing devices and technology invisible. To achieve this, a wide variety of smart devices must be deployed that cooperate with each other and fetch and carry information and services to individuals, often autonomously. And, of course, as additional devices are connected and more transactions take place, the infrastructure will need to be scalable, reliable and readily available.

This chapter aims to set out some of the key elements of this infrastructure, without too much jargon and with sufficient factual depth. If it becomes too dry and dull for the general reader, please skip through it and proceed to the third chapter.

To illustrate a path for infrastructure development, it helps to distinguish between different waves of pervasive connectivity. The current first wave includes the emergence of information access devices using a basic communication infrastructure, allowing the delivery of information at the point of need, through any device, from any source. In a second wave the devices will assume a more autonomic role. They will still use the same communications infrastructure but, in addition, sensors will observe human activities and help with specific tasks, at times without an overt request or command. This implies the ability of the network to recognise situations and contexts.

An early example of this second wave is given by the Microsoft "paperclip" assistant and spell checker, which pops up in Word and asks "do you need help with this?" The assistant still demands a steering decision from the user, whilst the spell checker operates automatically, correcting simple typing errors. A reader that has experienced frustration with such services may also wish to read this example as an indication of how difficult it is for a machine to derive valid hypotheses on the user's intentions from a set context. In a third wave, increased collaboration between services and devices would also be possible; thus, one system observes an action and reads the context, whilst a second fetches additional information and provides auxiliary services.

So *pervasive computing* will use the current network and communications infrastructure, accelerate its convergence and extend it to embrace sensory networks. It will provide additional "input" on the environment in which the users operate, and allow different devices to collaborate more autonomously than they do today.

2.2 Devices

A *pervasive computing* environment consists of smart interconnected devices. The spectrum of devices is broad and complicated. Ideally one would think of a device as roughly synonymous with the service it delivers: a washing machine washes clothes, how it does so and what processes it uses are hidden from view. However, if one thinks of a car we should be able to define it as a device that gives transport from A to B, but through decades of social usage this device has developed additional connotations of social status and individual identity that have little to do with safe or direct transport. In the same way the personal computer has become much more than the sum of the services it delivers, generating new patterns of behaviour and entertainment that make the computer apply to all sorts of situations.

For purposes of clarification it is possible to distinguish the following categories of device:

1. infrastructure – communications networks, sensors, actuators, RFID tags and readers, etc
2. access devices – that can be seen and touched, such as PDAs and laptops
3. embedded intelligence – controllers in washing machines, cars, cell phones, etc
4. symbolic keys – tangible “things” that are convenient representations of services in the real world. The credit card, for example, is essentially a tangible representation of a payment and credit service

An important effect of the trend towards *pervasive computing* is that the boundaries between the four categories of device will become blurred. Access device to the virtual world, such as the mobile phone may be used as symbolic keys giving access to real space, e. g. to an event or to public transport. More important, access to the virtual world will increasingly be provided by everyday objects with embedded intelligence and not by dedicated access devices. Placing some food on the kitchen table, for example, could activate a search for nutritional information and cooking recipes.

Whether future developments are dominated by a single “hub” that integrates the four device types described above, or by many different devices dedicated to specific functions, will be determined by culture and business competition as much as by technology. What is clear, however, is that devices operating in the background will grow in number and diversity. The number of mobile components per person will increase into the hundreds. Current estimates suggest the number of communicating data devices will grow from about 2 billion to 20 billion in five or six years, and could reach one trillion within another ten years or so. This applies to devices we may not normally consider to be communicating devices, such as toys, household appliances, automobiles and other machines. Some of these devices are tangible and others are embedded and essentially hidden in the infrastructure.

Radio Frequency Identification (RFID)
RFID is a method of entering data wirelessly into an information system, using devices called RFID tags or transponders. These tags are small objects, such as an adhesive sticker or button-sized device, which can be attached to, or incorporated into, a product. They contain antennas to enable them to receive and respond to radio frequency queries from an RFID transceiver. RFID tags are expected to replace bar codes, because they can be read and written automatically without sight contact, and store much more data.

Typical passive RFID tags cost approximately 25 cents (US$) today and there is the potential, using current manufacturing techniques, for the price to fall further, potentially to 18 cents (US$). After this, a further drop in price requires the evolution of the manufacturing technologies involved, specifically, in how the chips are bonded to the carrier material and antenna. The anticipated mass deployment of RFID will be heavily influenced by this decrease in their average price, and by the emergence of a broad industry standard consortium that agrees on a standard for RFID.

Passive RFID is just one variety of RFID technology. There are also "active RFID" tags, which are more expensive, bigger and have battery power. They are capable of transferring data faster over longer distances and with greater reliability. These tags are used, for example, to identify train carriages as they pass under a bridge, or to levy tolls on roads.

2.3 User interfaces

Sophisticated technology requires an interface allowing the user to interact with it in a convenient and practical manner. But if the computer of the future is essentially unobtrusive and invisible, how will this interaction take place? Over the past twenty years the computer interface has changed: from single applications on single screens to multiple applications in "windows", to the browser model with location-independent presentation, to portals with multiple websites integrated into a single view. These phases have all, however, occurred in the context of the "desktop metaphor" of human-computer interaction. For the vision of *pervasive computing* to become reality, more radical change in the interface between man and machine is required.

There is a clear trend toward communication through several modes at the same time, potentially combining voice, gesture and point and click control. In addition to understanding and analysing what is said, new systems are under development that can interpret gestures and facial expressions and enter into a conversation with the user. Research has revealed that the use of speech and gestures can speed up the communication with a computer five-fold compared with traditional menu-based interaction.

BlueSpace – the office of the future

Jamie Johns has his own office. When he walks into it the electronic name plaque on the outside wall switches from "Jamie is out" to "Jamie is in" and a signal light turns from blue to green. An electronic sensor, picking up signals from a chip embedded in his ID card, adjusts the desk chair, the table, even the temperature, according to his preferences. Urgent memos –"Jamie, I need you for a presentation in 5 minutes!"—flash on the wall or, perhaps, on the table top. Nearly every surface can operate like a computer screen. As he logs onto the computer he sees which members of a work team are available. If he does not want to be bothered he simply touches a screen, and the outside sign shifts to "Jamie is busy" and the outside signal light changes to red.

The overhead lights and technology controls are embedded in a movable structure that holds a projector with a 180-degree range to beam memos, spreadsheets or any other image onto the wall, table or floor. If someone without the proper electronic ID walks in, the display switches to a generic image, thus keeping sensitive information hidden from unauthorised eyes. Blue Space is a vision of how technology might solve a myriad of frustrations in daily working life, from uncomfortable temperatures, to elusive colleagues, to cramped meeting space. The project focuses on integrating architecture, furniture and technology and letting users control their environment.

2.4 Network infrastructure

To make new smart devices pervasive, networks without cables are needed, especially as it is essential to be able to provide users with access regardless of their physical location. So far, the market has been segmented into two categories: wide area networks (WAN) using cellular networks and wireless local area networks (WLAN) using an 802.11 type network. These implementations differ: WAN is specifically tailored to the limited capabilities of the network and devices, whereas WLAN takes the wired applications and PCs and puts them into a WLAN laptop with minimal changes. There is an emerging set of wireless technologies called Metropolitan Area Networks (MAN) that seek to close the gap between the cellular and WLAN worlds.

Different standards and capabilities need to be considered in a converged "anytime, anywhere" world supporting wireless broadband data capability. For personal area networks, Bluetooth operates for distances of approximately 30 feet, and WLAN provides distances up to hundreds of feet. Wireless metropolitan area networks are generally considered last-mile high-speed back-haul technologies to the home or office, and WAN, for example cellular or satellite networks, provides coverage over many miles.

To make *pervasive computing* work in a transparent and seamless way, it must be possible to go back and forth across these network types, something which is not straightforward today. For example, imagine you have a laptop computer, you still need to worry about connection to the appropriate network, and you may need Bluetooth to talk to another local device, such as a printer.

The nature of future networks will also change: they will become much more heterogeneous. Data, voice and multimedia carried over networks running the Internet Protocol (IP) are anticipated. The Internet Protocol is also expected to take over some of these standards. This leads in the direction of self-configuring networks. The wide variety of different applications will drive this as they seek to communicate with networks and each other without human intervention.

A wide range of applications could be enabled by the personal mobile hub: a new platform for deploying mobile applications. One example is Mobile Health (MHealth), which essentially acts as a bridge between a WAN and the local area devices that are in personal use. A blood pressure monitor that takes a reading, for example, transfers the data to a device that can transmit to a health care provider, which might check for an alert, refer the patient to a doctor or instruct the patient to check his or her blood pressure every three days. Access clearance must also be determined, for example, whether the doctor should have access to the patient's records. At the patient's next appointment, the doctor can see whether instructions to measure blood pressure have been carried out as directed and can discuss the appropriate course of action with the patient.

2.5 Standards and enabling incrementalism

In the future, communication networks will connect millions of diverse sensors, users, services and computer devices. Each one will be equipped with different processing and computing capabilities and will rely on the services provided by other participants within the network to complete its task. Infrastructure standards are necessary to ensure that these identities, or names, can be meaningful to all participants within the network.

We have discussed elsewhere the automated drive of devices to communicate and share information and, potentially, software components. For this to occur there must be some process by which different devices meet one another, rather like a handshake. Such technical handshakes already happen. What is required now is to enable these handshakes between more devices. A successful handshake depends on the following criteria: the two parties need to be able to find each other; there needs to be a common language in terms of frequency and coding; and the device must have a means of self-description and of understanding the description of the other device. This handshake has now gone beyond the simple expectation of finding another device that speaks the same language, it also implies the device can give a description of the services it offers. Finally there must be a form of higher level intelligence or overview to be able to determine that the two individual services can form a synergy. It is of course assumed that the description would contain security and privacy rules to avoid secure devices collaborating with insecure ones.

This handshake is an example of the type of enabling technology that will be required to exploit the combinatory power of so many individual devices. The final vision is a self-organising process, with devices seeking each other out, combining existing services in order to deliver new ones. The result would be an evolution of services and of the structures used for providing them.

So far it has been assumed that there is a correlation between the pace of new technology uptake and the development of industry standards. Whether it is better for these standards to be controlled by independent bodies, major companies able to invest and expand the market or new, hungry spin-off businesses is irrelevant. Nor would it be interesting to make a great appeal for the creation of industry standards, necessary though they may be. One approach to standards is to consider enabling incrementalism: the incremental element recognises that the initial steps of *pervasive computing* will be partial and occur in stages. The challenge for researchers and developers is to identify services and applications that are robust enough to deserve investment, without the requirement of a future where *pervasive computing* is fully implemented within a more general context. An example would be in logistics, where there is a vision of an integrated industry supply chain, where each link of the chain has a full knowledge of the sales patterns, stocking, spare parts and delivery times of all others. This may be a logistics nirvana, but the reality is that each step of the way towards this vision is being keenly fought for in terms of a cost and benefit approach.

This incrementalism also makes sense, given that new services and applications only emerge once the requisite level of connectivity has been achieved. In the prevailing business climate, investment is going to be made for services that can be concretely defined today and not for the "intelligent fridges" of the future. In contrast to this, there is also a need to prepare the infrastructure today that can expand once new opportunities emerge. These enabling elements will play a critical part in allowing the benefits of *pervasive computing* to be experienced by users.

2.6 Network traffic

Sensors and other devices produce a large amount of data at the edges of the network. This data is sent to a server, which somehow has to deal with it. This can cause grave data aggregation problems. Depending on how they are programmed, sensors can be very greedy collectors of data, especially if they are tasked to collect "real time". Like digital hoovers, they suck up data and pass it onto more intelligent machines that can aggregate and sort it to create information. In order to illustrate the trend, take the flow of data which is transmitted from the edge of a network towards the centre. With a sensor to server configuration, the ratio of data flow in both directions will be in the order of 100:1. This inverts the current client/server paradigm, where the so-called client on the edge of the network sends requests to a more central server that is working in parallel with many other clients. The data transmission ratio in this existing configuration of client to server is 1:50. This means that, to manage the aggregation of data from sensors, a different type of network infrastructure will be required. The aim will be to try and filter and aggregate the data at the "edge" of the network, closer to where the data is collected, leaving servers to perform more important tasks.

Section 3: Unpredictability, Collaboration and Optimisation

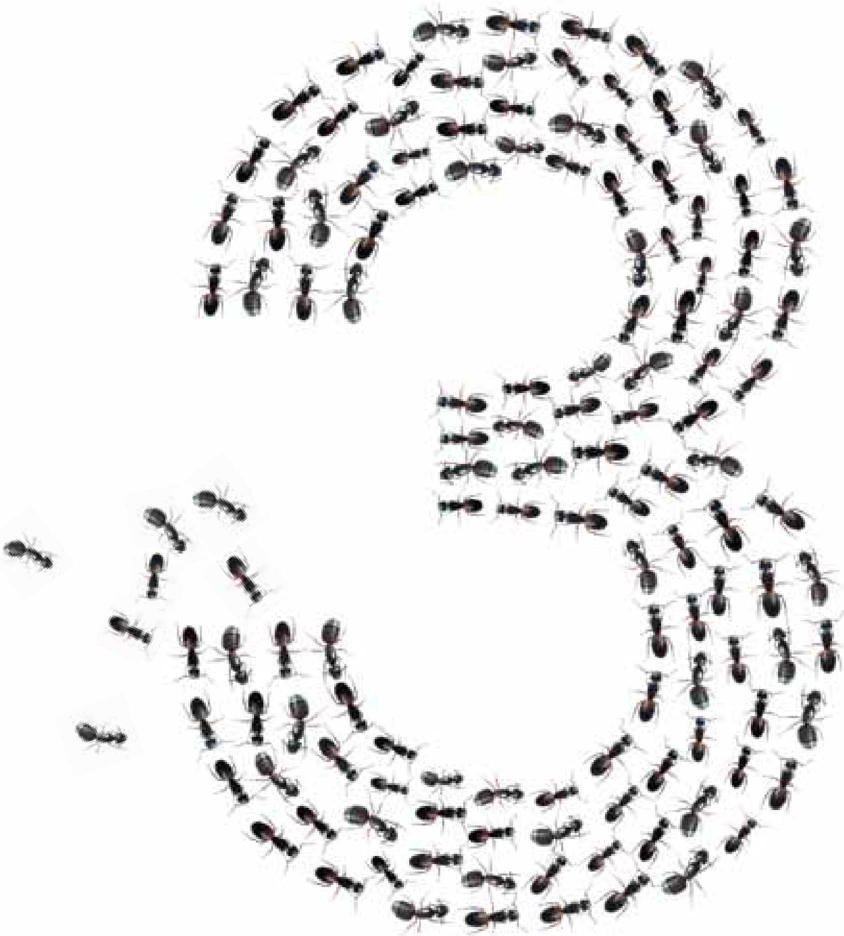

3.1 System fragility and criticality

When explaining what might be new about *pervasive computing*, experts often employ terms such as “self organisation” or “system equilibrium”, stressing the inter-connectedness and fragile dependencies within natural and man-made networks. Rather than hierarchical systems, stacked up like a house of cards, the imagery is of a swirling cluster of nodes interacting with one another, dynamically changing without external intervention. Unlike the analogy with the stable and robust self-organising ant colony, however, these systems also suggest that “critical events” may occur.

A simple example is Bak`s pile of sand: like a child on the beach, steadily pour sand onto a flat surface and it will form a pile with a regular form. Keep pouring and sooner or later one side or another or all of the peak will slip away. The pile of sand edges towards a “critical state”, one where the form becomes fragile and the future structure unpredictable. Keep making new piles of sand and each time the response is potentially different. Occasionally major shifts occur, such as the whole cone collapsing, but most of the time much smaller adjustments are made.

This line of thinking is very seductive. It is possible to see complex systems edging towards criticality everywhere, from stock market crashes to traffic jams, to electricity “blackouts”. The relevance of this approach to *pervasive computing* lies in the perceived increase in the complexity of networked systems that will come into operation and their potential to edge towards an irreversible criticality. Of course there are many examples of complex systems that successfully maintain their equilibrium through time and indeed where intervention can be as harmful as helpful. However, to avoid catastrophic outcomes and reap the benefits of resource optimisation and improved system performance, there is a need to increase our awareness and understanding of the limitations and adaptations we may face in managing and using systems of the complexity envisaged with *pervasive computing*.

Electricity grids are one contemporary example of the practical issues this increase in networking can entail. Integrating electricity grids makes sense because electricity cannot be stored, so the provider only wants to supply the amount needed to satisfy demand at any given time. A larger and more diversified network can do this effectively as it can manage the peaks and troughs throughout the system. The practical element, however, complicates this theoretical approach. Grid systems tend to have been built and updated incrementally, according to different standards and performance expectations. Put them together and the risk of unpredicted events occurring increases. Minor events, such as a tree falling on a power cable, if not managed adequately, can lead to power surges or reductions cascading through the system that result in a widespread system crash.

One of the difficulties with these systems is that most of the time they function extremely well within defined thresholds and with reasonably predictable outcomes. However, when they step outside of these thresholds they can reach a state of criticality where the outcomes are unpredictable and potentially catastrophic.

A river system is a relatively straightforward network compared to decentralised networks that allow communications in two directions. With the heavy rainfall in Central Europe in August 2005, however, small mountain streams had a throughput of water which was beyond the normal thresholds. The water was able to undermine the previously stable banks of streams causing trees and mud to fall into the moving torrent, which was also abnormally loaded with glacier debris dislodged through melted ice, due to warmer than average conditions. At times the debris coming downstream would get caught on rocks forming an artificial dam and re-directing the flow of the stream for a short period, leading to further erosion and potential flooding. Finally a stream would reach the valley floor with unpredicted force, risking a further wave of unanticipated damage.

Such system-related events create questions as to who is liable and responsible for specific damage. In the case of the electricity grid it is possible to trace the physics of the power surges back to specific events and assess the grounds on which appropriate action was not taken to avoid a blackout. With more complex interactive systems, however, this is not always the case, and if catastrophic events occur quasi-automatically through the interaction of different systems, who can determine the boundaries of ownership and identify the system provider, define their responsibilities and prove liability?

3.2 Designing systems and unpredictability

The type of connectivity described in the first section of this paper can lead to a level of hyper-networking where the resulting system no longer has a designer or provider, and there are so many different interactions between different parts of the system that it is not possible to always predict the outcomes.

Traditionally a good systems engineer wants to define the boundaries of the system he or she is designing and test all the possible states the system can be in. For example, that it will achieve "graceful degradation", meaning less important parts of the system will shut down first in response to a change in conditions and that measures will be taken to avoid "harm to life", just as when a signal in a rail transport system fails, all the trains stop dead until the signal is fixed or an alternative routing plan is confirmed.
This elegant and intelligent approach to

designing systems cannot stop all failures, because systems can only be tested for their intended uses, having too many possible states to be tested exhaustively. If a system is used in a different context, it may give different outcomes from those intended, which may create new and unforeseen interactions with other systems. Every research laboratory has a story of how a new system design, after rigorous testing, was placed in jeopardy through the cleaner switching off the power or the security guard accidentally jamming a signal with their short-wave radio.

A systems engineer learns and practises rules of design, but he or she is fighting against the pace and scale of technical developments. There is pressure for the designer to reuse "modular" pieces of software written in a different context to speed up the development of a new system. If each element is adequately tested this should not be a problem.

Integrating larger systems, however, is different. Put existing systems together and the number of system states increases exponentially with the number of connected systems. The bigger the systems being integrated, the more difficult it is to deal with all their possible states. The logical and controlled environment of the system engineer reaches a certain limit. A new level of unpredictability enters about how systems will interact. The emphasis therefore moves from trying to validate a system to trying to manage the uncertainty of its potential outcomes. This is a necessary consequence of the idea of self-organising systems with emergent properties.

3.3 Who or what is the circuit breaker?

Envisage a hyper-networked society, where systems and objects are constantly trying to make contact with other systems in order to share software components on a common platform or build a specific application or establish cooperation to provide a new service or share data. One feature of this type of integration is that it is very difficult to go back or unwind what has already been done. And humans, who often play the role of "circuit breakers" today, do not have the same role within the "control loop". This would present a new challenge of understanding how a hyper-networked, *pervasive computing* environment will operate. How do different parts of the natural and IT worlds connect in a non-linear way? And how should one manage a situation where one pushes a particular button to achieve a desired effect and it influences a dozen other variables that negate the desired outcome and initiate further unintended processes?

Despite many years of debate on artificial intelligence and the potential of computers to act intelligently, it *appears* there is an issue as to how broader communities will use or avoid systems that appear to operate in a non-deterministic and intelligent fashion. Appear is in italics because there is the potential to assign intelligence to systems because of their speed of calculation, their combination of different sources of information or because they are simply too large to comprehend, none of which require the system to be non-deterministic or unpredictable. There is a tendency towards wanting to assign intelligence to a system without breaking the expectation that it will be both stable and trustworthy. If unpredictability, as well as intelligence, is an emergent property of complex systems, the question remains as to how to respond when critical or catastrophic events occur. Is it best to call them unavoidable "freak events" and continue to trust the systems that generate them or will this deeply influence the acceptance of this type of background collaborative computing?

3.4 Optimisation

Optimisation applies the discipline of mathematics to create a model of real life situations in order to find the best solutions to a given problem or to simulate changes in a system by adjusting the required inputs and outputs. Optimisation is undergoing a revolution through improvements in processing power, data availability and mathematical or stochastic methodology. Importantly optimisation is, in this context, not about finding short-term "peak performance", but rather finding the best solution based on the weighting of the initial parameters and the goal that is to be achieved within the computer model.

Connecting objects to digital information provides unprecedented access to data at a time when progress in data storage and analysis is accelerating. Information will increasingly be available instantaneously for assessment against large statistical banks of similar data. Further use of optimisation systems appears almost inevitable. There is not only a powerful economic logic to such an approach, it also fits our current focus on the efficient management of public and private services. Twenty-five years ago, manufacturers would predict demand once a year, ten years ago this was done on a quarterly basis, today it is once a week. The challenge is that optimisation calculated over shorter periods tends to lead to greater reactions to short-term variables and therefore a greater volatility in the results, something that is "smoothed" out when viewed over a longer period. There is therefore a need for us to understand the choices implicit in optimising a specific service and ensure the different stakeholders involved understand the consequences of the perceived improvements.

There are also several important social

aspects to optimisation, which link into game theory, economic performance and political decision making. A financial investor may, for example, argue that a public company must do nothing but focus on optimising short-term performance because they see this company as one part of a broader financial portfolio that they themselves are seeking to "optimise" for maximum financial return at an acceptable level of risk. On the other hand, the company executive may see their role as to optimise for the longer term, generating employee and customer loyalty that will generate good and sustainable profitability. This example does not pass judgement on either party, it simply shows the collaborative and competitive nature of optimisation within a complex system.

As the ability to model reality and systematically simulate and optimise outcomes improves, we must take certain issues into account:

- What are the boundaries to optimisation: which variables are being measured and are some factors excluded because they are difficult to quantify? What is the goal of optimisation? Is it a sustainable outcome or short-term peak performance?
- What influence can the optimisation have on factors external to the model? Have the potential stakeholder responses to optimisation been taken into consideration, e. g. what will happen if competitors use the same or similar optimisation models?

Resource Management

The flow of information helps to improve estimations of what resources will be required within any process or system. Although this may significantly improve resource management from the point of view of the investor or generator of inputs to the process or system, it can also create surprising outputs for different parties. In order to run a system close to peak performance, it is necessary to take it close to the point of breakdown. A Formula One car does this to maximise speed and responsiveness over the race circuit. The New York road system does it to minimise tax payers' capital investment. It is also possible to optimise a system to the goal of stability and reliability, such as the Zurich tram system, where the system uses relatively old vehicles and infrastructure but focuses on maintaining frequency, with a tram stop always in reach and with very few breakdowns or changes to the schedule.

Several factors emerge from *pervasive computing* in terms of optimising systems. It further increases our potential to design systems either for sustainability or for the shorter term. It may also help integrate other factors into the optimising process that are currently disregarded. If a manufacturer doubles its output of furniture by dropping the price by 30% and reducing the life-span of the furniture by half, he may significantly increase profit. But this "just good enough" product will also use more natural resources and lead to more waste and pollution. Such factors are external to the optimisation calculation the manufacturer needs to make today. *Pervasive computing* may also allow costs external to the model calculation. It may also enable approximation of the long-term total cost of projects that could be calculated and adjusted over the life of the project.

3.5 What will the increase in data and analysis do to us?

It is very possible that optimising systems will generate sub-optimal experiences for users. Try purchasing a car with a reduced specification, without leather seats for example, and you may be told it will take six weeks longer to deliver. This is sometimes done deliberately to influence choice, but is also because the design of the system was achieved against different criteria than those that will satisfy the individual user.

This flood of data can also be directed at the individual. The tendency to see the mind as a super-computer and describe the body as a machine is increasing. Athletes are already measuring and recording the functions of the body to understand how they may influence it to improve performance. In the future, such tools will be available to all. The influence of a heavy meal, and perhaps too much wine, will become apparent the following morning through the read out of a machine indicating a reduction in water content in the body and a change in the metabolic rate. People will be able to compare statistics with friends, take actions to improve physical attributes and then measure the results.

There is also scope for saving lives by treating and predicting health problems before they become life threatening. It is assumed such data will remain confidential, nevertheless difficult ethical issues will arise. If the effects of specific patterns of behaviour can be followed through time, such as poor nutrition, stress, lack of physical action or dangerous sports activities, and an individual recognises but refuses to change their pattern of behaviour to ameliorate the symptoms, should the state or insurance company pay for their medical care? Will the "pay-per-risk" business model offered by some car insurance companies be transferred to the health insurance market? And what would be the optimisation goal of a health system based on real-time personal health monitoring?

Mediated travel

It is family vacation time and the first adventure for the four members of the Studer family is a trip to the Italian Alps. To save time, the journey planner has selected to use a car train rather than drive over one of the mountain passes. The journey will last about twenty minutes. During that time there is the opportunity to have an update on progress. First, there is the question of driving performance. The on-board computer benchmarks performance against that of the notional "perfect driver". It beams examples of how the turn into the bends is consistently too quick onto the windscreen, leading to over-braking and an estimated 2% increase in fuel consumption, due to the need to re-accelerate. It also shows how safe braking distance to the car in front is not being respected. As one can imagine, the other passengers are not impressed.

One of the children points out that the car already adjusts the braking, gear change and steering, so the driver should be able to do better than this. Without comment the driver turns swiftly to an insurance update. A pay-per-journey package has been selected as few journeys are made abroad. Fortunately, the update confirms that the driving performance remains within the agreed price cover, but indicates it is snowing in the next valley requiring an average reduction in speed of 20km per hour to maintain the same insurance rating. It looks like the arrival time shall be later than planned.

Meanwhile the children, who stubbornly refused to look at the view, have taken photos of some of the mountains and are using a combination of satellite pictures and editing tools to redesign the landscape, making a perfect 500m long snowboard park, complete with half-pipe and jumps. To satisfy their parents they also send information to the windscreen on the real landscape they are remodelling, indicating the old mule trails and the possible places to eat in the area. They are just quick enough to delete by voice activation an update to their model, which has been beamed in from a friend at home. The update indicated the friend's research on the hidden military complexes within the landscape and his playful plan for how to destroy them.

Finally the end of the train tunnel comes in sight and the flurries of snow sweep around the car obscuring the view. It is good to escape the mediated landscape the kids have created, but now time to turn attention once again to the driving, with renewed determination to improve the driving performance on the winding road. After all, no machine is going to get the better of a human.

Section 4: Underlying Social Trends Affected by Pervasive Computing

This chapter records some of the observations made during the workshops by the expert group on how to balance the social impact of pervasive computing with its potential perceived benefits. Many of the factors under discussion were perceived by the expert group to be neither intrinsically good nor bad. It was also felt that a high level of social awareness would play an important part in avoiding highly volatile reactions to potential social changes and to ensure intervention could occur early to prevent adverse developments.

1. Inside and outside the systems

In the past it has been relatively easy to distinguish the machine from the human, but will this continue if humans become another class of object within the digital environment? Of course, if machines are to be helpful in anticipating human needs, they must know what individuals are doing, interpret the context in which their actions are occurring and even predict what a person will do next. For some this is a Faustian signing away of human freedom, for others it is a question of adapting worn-out habits.

If there is a breakdown of the subject-object split between man and machine, it is also because the machines are going to get on with things without us. Will the machines perform their delegated work so well that humans are simply superfluous to requirements? This is rather like the Sorcerer's Apprentice, only this time, hopefully, the objects remain under control. The objects proceed, dispassionate and uninterested in thoughts and feelings, completely focused on fulfilling their allotted tasks, answering the perceived needs of the users, regardless of the struggle the user may have in expressing what these needs might be.

2. Virtual merging of our social, family and working roles

Is *pervasive computing* an extension of the "anywhere, anytime" culture, which many have experienced through existing mobile communications and computing technologies? This has forced new flexible boundaries between the different spheres of work, home and leisure, leading for some to a sense of increased stress and for others to greater empowerment.

3. Learning new tricks and forgetting old ones

Pervasive computing presents a move away from interfacing with computational devices via menus and screens, towards networked computational devices that are present throughout the environment. This implies a mutual adaptation between the technology and its user. This computing environment will be unable to perfectly adapt to explicit requests or to correctly read the context or user intentions. New habits will therefore be acquired or "tricks" to let the appropriate interface know what is desired, or even to cheat it in order to avoid undesired reactions. The systems will build user models, and the users will build their own approach to deal with them. The unpredictability and intended unobtrusiveness of the systems will make this a harder task for the user than before.

4. Adjusting our problem solving skills

Is there a threat of deskilling? General advances in understanding and manipulating complex technology may lead to a reduction in the ability to solve problems that may appear relatively simple today. This is similar to the frustration an avid user of in-car navigation (GPS) feels when he or she must use a map again. It can be argued that this concern has always been expressed by the older generation to the younger.

A sixty-year-old today may remember hours of fantasy play as a child using nothing but wooden bricks and imagination. For them, their grandchild simply "consumes" television and computer games. For the child, however, the computer game is a fascinating challenge: what must the player do next? And how can the player develop the right strategy to fully define and reach the objective of the game?

That three teenagers may need to send multiple SMS messages to each other to move from a vague agreement to meet, to an actual time and place for the meeting, may appear highly inefficient to an older generation who would have had no choice but to agree the details directly and simply keep to the arrangement made. But this indicates that the technology enables new forms of social interaction and role playing that do not always imply a direct improvement in efficiency or use of time. *Pervasive computing* will therefore, at this quite basic cultural level, influence which new skills we learn and which ones become obsolescent.

5. Altering social interaction

A related issue is that of social interaction skills. Think of a group of business executives enjoying a coffee break in a meeting by all individually phoning their offices or friends, all isolated from one another and their immediate surroundings. With *pervasive computing* public transport tickets will be purchased automatically and billed electronically, simply by passing through the appropriate gate. Asking directions will not be necessary as people will be guided by a hand-held navigation system. And whatever news, be it financial, weather or sport will be readily available, again tailored to personal requirements.

In such a world, will passing the time of day with a stranger standing in a queue still be a possibility? And will people be too busy interacting with another stranger via a virtual chat room through a mobile hub to really care?

6. Digital life maps

Privacy becomes an issue when data gets organised and attributed to individuals. Although significant digital traces are already left through mobile phones, credit cards and customer card usage, this does not seem to overly concern us for at least three reasons: first a belief that such data simply gets lost in the chaotic noise of the digital information sphere. Second, that any individual bit of data is not sufficient to reveal exactly what a person is doing, it remains a very faint trace of our historical activities. Third, the potentially misguided modesty, that asks "who could possibly be interested in what I do?"

If humans become a class of object around which their activities cluster in the form of digital data traces, there is the potential that a digital representation is formed of who that individual is that could be extremely rich in information. The aggregation of the information comes to represent a digital life map and the individual becomes essentially readable. A court of law, for example, may treat this digital representation as a much more accurate witness than any single person, especially if they are motivated to lie.

Of course, there is a distinction between being readable and being read. Is there some way that this digital representation could itself be a private information cluster? As soon as the individual is able to choose to whom he or she passes over their digital life map, a different feeling may emerge about sharing it. And, if there is a desire for systems to be able to anticipate needs, they need to be able to generate a rather accurate description of past individual actions and what the individual might do next.

7. Evasive technologies

Do humans have a "get-out" clause that allows them in some way to step outside of these interactive environments? Our experts felt there would be a demand for such a possibility. There may also be the potential for evasive technologies which would deliberately disrupt the smooth flow of information and digital traces. The creation of geographical areas in which electronic devices are not allowed would be similar to the creation of natural wildlife parks, where animals and vegetation are left to develop an ecosystem, to some extent free of further intervention. There may also be a demand for protective or counteracting technologies that interrupt wireless networks or aggressively deny undesired services. Such services may play an important role in increasing a feeling of self-determination and control for individuals, even if they are not widely used.

8. Dynamic pricing

Payment for services can take on many more forms. The potential to adjust pricing according to supply and demand will increase. Even durable goods, for example a television set, may change in price according to the relevant companies' global stocking position. Products where supply is constrained may use the type of optimising programmes employed by airlines today: as seats on a particular flight are sold, the price for remaining seats is adjusted. Other companies may seek to exploit peaks in demand, such as fast food companies charging more for their goods over the lunch break, or a supermarket behaving like a stock market. A soft drinks machine which has access to the weather forecast may automatically stop selling drinks when it senses the temperature rising, because it can sell the same items for a higher price when it is hot.

Intuitively, these examples feel artificial. The experience of buying petrol, for example, is hardly improved by discovering as one leaves the station, that the price has fallen 1%. However, the advantages to companies in managing their supply and demand cycle could be sufficient for them to seek to prove to the customer that, over the long term, average prices can be reduced, even if the individual feels that prices always seem to be higher. Economic theory states that the more efficient markets are, the more wealth is being created, so there is a theoretical case for dynamic pricing. However, consumers may react ”irrationally“ and strictly reject dynamic pricing, as it demands substantially more ”attention“ when shopping and leads to decreased transparency.

Other payment systems may also be extended, such as pay-per-risk for insurance or pay-per-use in car leasing. The price of insurance for a car journey, for example, may be calculated by a set of factors including the speed of driving, type of road and weather conditions, volume of traffic and time of day. Such possibilities may have wide-ranging effects on industries, such as insurance, that currently set their pricing on large sets of statistical data and the mutuality principle.

9. Changing health management

The potential to monitor the body and regulate its activities will be radically extended in the next fifteen years, and not only in hospitals. The aim will be to prevent illness and maintain health in order to avoid costly medical interventions and care of the long-term sick. Arguably at the heart of this process will be the availability of real-time information on specific individuals. The causes of illness will be much more apparent, particularly when they relate to lifestyle choices, such as over consumption of alcohol and food. The funding of medical care will probably become closely linked to individual lifestyle choices.

The ski resort

A ski resort lends itself well to pervasive computing. There are a lot of optimisation problems, such as how to queue most effectively, especially for the long cable car trip up the mountain. There are issues of tracking people and equipment for safety, payment and to avoid theft. Finally there is also a desire to hide all this computing because people go to ski in order to experience a "natural" environment, even if it is within the infrastructure of mass tourism.

In a pervasive computing ski resort equipment will be "tagged" with the individual's identity. It will check with its owner's other processors (mobile phone, credit card) before unlocking the bindings, and sound an alarm if the boot is not applying full pressure to the ski due to a fault or snow stuck to the bottom of the boot. According to the type of ticket bought, the equipment will give access to the services available, automatically opening the turnstiles and giving a personal greeting in the morning as well as showing the number of the reserved cable car for the journey up the mountain. At the top of the cable car there may be a group of teenagers in the snowboarding park reviewing their jumps on a giant screen. It will replay their previous descents with a shadow performing the best line and speed. Of course, the shadow never falls over.

Another board could be flashing red dots against the mountain. On approach, two dots expand and text appears indicating the dots are your two children, out with their ski teacher. Two other dots are flashing a violent green indicating that someone is leaving the official off-piste area in search of fresh snow. The identity check indicates they have no avalanche training and have not cleared the trip with the authorities or their insurer. A ski guide is already on their way to indicate the dangers. There is no amplified music. Special ear implants are available and a wireless system allows you to choose what type of music you want from Mahler to Dolly Parton, Eminem to Kool and the Gang. It is time to ski.

In the evening, on returning to the hotel, the data house system transfers information collected and stored in the ski boot through the day. Lying in a steaming hot bath it is possible to review the day. Data has been collected through the sole of the boot indicating that dehydration is an issue, due to the dry cold and that there is an imbalance in leg strength that is affecting the ability to ski. It makes some suggestions about exercise to correct this. It also correlates this information with the chip monitoring long-term blood pressure and indicates that the altitude is having a good effect. It must be time for a fondue.

Section 5: Social Impact of Pervasive Computing and Privacy

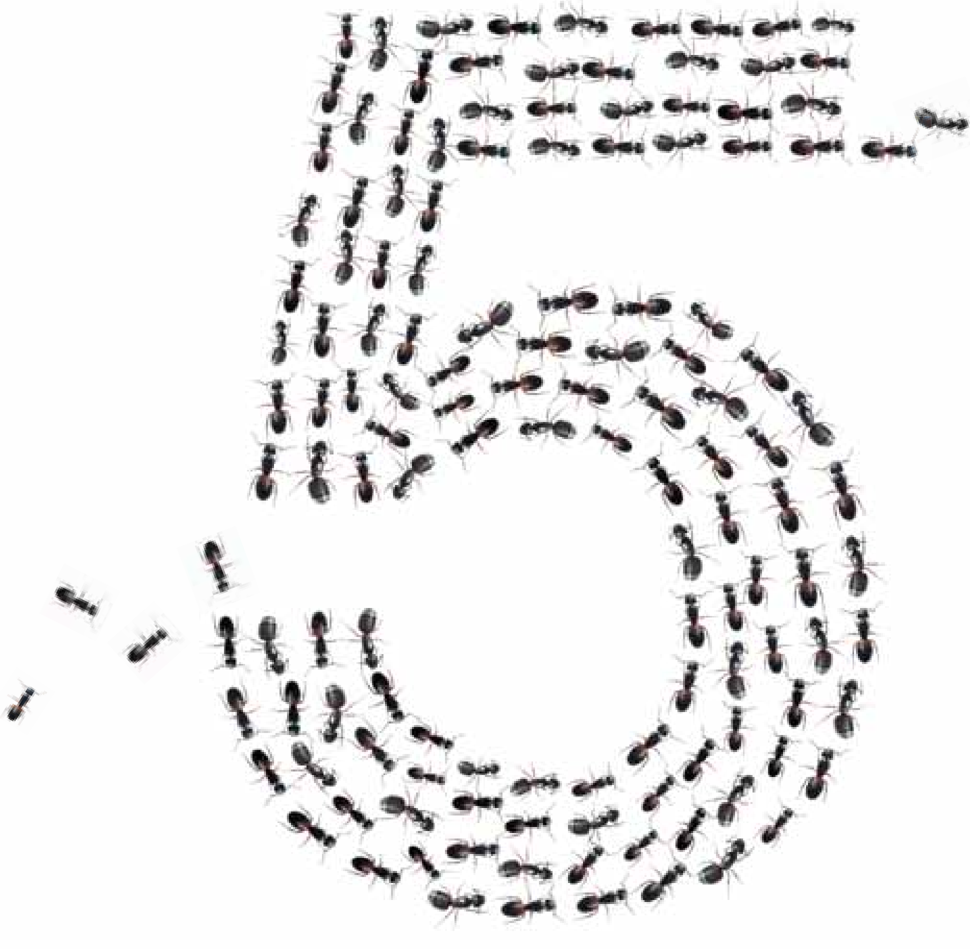

5.1 Introduction

Is loss of privacy the price that must be paid for the benefits of *pervasive computing*? If machines are to understand an individual's needs and act on their behalf, they have to know what that individual is doing and store and analyse what he or she has done in the past. The more sensor networks, bio readers and processing units that synchronise with mobile phones and other devices, the more digital traces are left, both intentionally and unintentionally. Mantras such as "any data, any time, anywhere" imply that the data can also be about "anyone", raising the question of whether these digital traces could be collected and organised to create a digital representation of an individual's actions? And if it were possible to create something like a "digital life map" of an individual, just how comprehensive would this life map be and who would be able to view it?

These digital traces would be more than a record of past transactions and locations visited. The actual digital traces left can be correlated with generic statistical data to create predictions about future patterns of behaviour. So a digital life map could stretch not only from the present into the past, but also from that past into the future, like playing a game of tennis or chess, where the opponent always seems to know what one is going to do next. If someone is thinking of buying a new car, perhaps a detailed offer is already lying in their post tray, not through an act of mind reading, but through extrapolation from existing information: two visits to the garage in the last month, a car more than three years old and a third child recently born. And perhaps the most troubling thought with this example is that in many cases the prediction may be correct.

In the early series of Star Trek, Captain Kirk was often faced with the dilemma of how to escape an all powerful computer that could always anticipate his next move. His escape route would take one of two forms: faced with the choice between saving his ship and all the crew against the life of the local ambassador's son, he would always make the ethical choice of the one over the many, instantly destroying the computer that could not compute this form of non-utilitarian thinking. The second form of escape was to present a paradox based on self-reference, "what I am thinking cannot be thought", which would freeze the all powerful machine because it could not compute the contradiction.

There is no reason to assume that anyone will be faced with an all-powerful evil machine wanting to destroy their space ship in the near future, but perhaps Captain Kirk is still beaming a message: one rich quality of the human approach to life is based on the ability to make decisions that go beyond the boundaries of mechanistic calculation and self-referencing paradoxes. Privacy is bound up with issues of freedom, identity and ethical and political choice. And the question therefore arises, if the individual stops being a sovereign subject within a world of objects and loses both the physical and psychological sense of privacy, what will happen to the notion of privacy?

The topic of privacy can appear overwhelming, acting as a magnet for concerns with the pace and scope of technological change. The temptation is therefore to either confine the scope of the privacy discussion to that of data protection, itself a complex and difficult subject, or seek to dismiss the issue: the notion of privacy simply adapts itself over time, it is a thing of habit which is traded for personal benefit in the form of greater security, cheaper prices or other incentives. The danger with these approaches is that the problem gets explained away, but without removing the underlying concerns that are being expressed. To avoid this type of dismissal is especially important when viewing future risks from a social perspective, as how an issue is perceived can have a greater impact than the actual hazard itself.

This perception of risk is of course highly influenced by the media, which is looking for powerful stories and imagery to convey difficult subjects and at times to generate sensational effect. In the case of *pervasive computing* these depictions may focus on fears over loss of privacy, described in terms of a *pervasive computing* "end-game", a nightmare scenario dominated by images of Big Brother. So, while the introduction of *pervasive computing* type technologies is incremental and uncertain, the perception develops of an inevitable journey towards a society where everything and everyone is monitored, with or without their consent.

The approach used here is to take the perception of the loss of privacy seriously, that such concerns express legitimate fears about a potential future with a negative view of *pervasive computing*, even if this is only one possible outcome. Of course, if this end-game were to occur, it would be the product of a cumulative set of choices made by individuals, small privacy trade-offs and collective decisions made in response to security concerns. This approach, therefore, sets out a strong version of the *pervasive computing* "end-game" and explores why privacy becomes so central in the debate on new technologies.

5.2 Shifting boundaries of privacy

Privacy lies on a fault line between maintaining individual autonomy and the communal or state demand that individuals be accountable for their actions. This image of the fault line is used because it expresses a dynamic aspect to privacy that shifts through time and across cultures and often acts as a point of tension and release for debates on the balance between personal rights and social duties. The fault line can undergo seismic shifts, as seen with the Patriots Act in the USA. In this instance, driven by external threat, the balance has tilted between the security needs of society and the privacy rights of individuals.

Many trade-offs and negotiations occur along the privacy fault-line. It seems there is broad, but not universal, consent to give up small elements of privacy in return for minor incentives, such as a retailer that offers a customer store card, including discount prices, with the explicit intention of gaining detailed information on customer consumption patterns. Consent to this minor trade-off is based on many factors, for example that there is freedom of choice as to whether one accepts the offer and under what conditions the information is used. There is also an expectation of a form of symmetry between giving up an element of privacy and getting something in return and that the increase in personal accountability is not greater than the perceived benefit.

Within legal limits, the private sphere is that which remains non-accountable, a place and time where the individual is not accountable to anybody for his or her actions. Some would argue this potential to experiment and explore different aspects of ourselves within the private sphere is vital to learning to make good collective decisions in public. This includes the right to take on other identities and play out roles quite distinct from those that operate in public.

On each of these counts it appears *pervasive computing* challenges how the privacy fault-line is configured. If people leave digital traces as a matter of course, it is difficult to give consent to them being kept and there is a major risk of an asymmetry arising between what is given up and what is received in return. So, it appears that these digital traces will increase accountability and reduce privacy. Yet it must also be recognised that, because the notion of privacy changes over time, the privacy fault-line may well reconfigure itself so that society will not actually be overly concerned with the new state of affairs.

5.3 *Pervasive computing* "end game"

A journalist often negotiates with the person they want to interview to establish the terms and conditions according to which the material will be used. It can be an "off-the-record" briefing that will remain non-attributable or "on-the-record" and therefore generating material for public use. This could be one way to understand privacy within a pervasive computing environment. The default switch is reversed: previously everything was private until made public, in this new environment, everything would be public until the switch would be turned. Everything is "on-the-record", all the time.

Of course if the prevailing context is one where a digital representation or record is always available, then the act of removing oneself from this public arena, of switching to privacy, can also be read as an evasive act: is there something to hide?

This is an over-simplification and it is important not to slip into describing such an "on-the-record" environment as one that is actively managed by a state, any organisations or indeed individuals. At the most simple level, the collection, storage and analysis of data is simply an integral part of getting machines to perform background tasks without human intervention. After all, because a message can be read does not mean to say somebody is reading it. It does not imply that all the information being collected is for the benefit of others, nor that there may not be many benefits for the individual in belonging to such a collective information sphere.

How will an "on-the-record" environment change things for the individual? He or she shall certainly be more accountable. Infringement of minor laws, such as speeding or not paying for a train ticket can be discovered and penalised, without exception. This does not mean an end to crime, but certainly the forms it takes will change. Tracking cars or other possessions will become much easier through embedded processors capable of signalling their location. As discussed earlier, car insurance may be based on actual and individual driving performance rather than through a premium calculated for a particular risk group defined by age, sex or nationality.

Through the second half of the last century the car became a complex symbol of freedom and an example of a quasi-private space that occupies a very public environment. Many continue to prefer to commute by car, despite delays and risk of accident, than travel by public transit systems. And there can be few better examples of technology being used by individuals to express their personality and status. It is perhaps impossible to surmise what effect a new type of accountability may have on a person's relationship with the car. Pay as you drive insurance and instant speeding and parking fines will change habits, as will semi-automatic driving systems that regulate overall speeds and traffic flows. Will our communities simply adjust to these changes and find new toys through which to express identity? Or will individuals rebel, as some small groups are opposing the increasing use of speed cameras in Great Britain, finding in this execution of the law an illegitimate imposition on their personal liberty?

At first glance this appears to be an area where the language of the "Big Brother" state seems to fit. Among regulators and legislators, however, there appears to be a sense that the collective social framework depends upon a level of individual freedom. There is a reluctance to see this private sphere become accountable because it is an opportunity to test boundaries and make risk assessments in order to have the skills and awareness to function within and for the collective.

What is in it for the individual? Is there an asymmetry between what companies and governments get out of collecting data on individuals and what the individuals get in return? One argument is that too many of the benefits of *pervasive computing* sit with large companies and government as it is currently envisaged and that this will lead to individuals fighting against its implementation, often on the issue of privacy. This could be a question of perception. Any company would argue the data they collect feeds back into better tailoring a product for the individual which will lead to better price and performance for the customer.

Most data traces, even without including a name or identification number, are potentially related to the identity of the person who left them. Data stored without an explicit reference to our identity can later, by statistical analysis, very often be related to a set of individuals. The accuracy of such retrospective analyses rises with the amount of available data. For this reason, acting anonymously may no longer be a sufficient condition for feeling private, free and unobserved.

Another issue is to what degree can one really give consent to the collection of data within a *pervasive computing* environment? The collection, storage, aggregation and analysis of information are going on everywhere and all the time. And this also raises the question of how people will be aware of data being collected? If it is going on all the time, is it possible to change it? Perhaps inevitably, the protection of privacy at this level belongs to one of the small daily battles which people become resigned to losing: niggling and frustrating but not high enough on the personal agenda to pursue further.

5.4 Will *pervasive computing* really make such a difference?

Will the quantitative growth of digital traces in various systems lead to a qualitative change in our environment? The expert group differed in their views on this topic in two ways: how big the changes in the environment would be and how strongly society would react to these changes.

Today, a trace is left behind each time a digital transaction is made, whether it is with a credit card, telephone or e-mail. In most cases, the individual chooses to trust the credit card company or phone operator not to share any personal information and indeed to actively protect it from theft. And in many cases companies are legally bound to store this data for a given period. The use of anonymous pre-paid phone cards by terrorists to arrange their activities and the extensive networks of websites maintained by extreme groups has strengthened the regulators' resolve over the past five years to ensure digital traces remain traceable over time.

Put these different digital traces together and you would already get a rich image of a person's location and activities. That many individuals feel sanguine regarding this fact indicates the delicate balance that prevails when dealing with privacy issues. Certainly the habit acquired incrementally of leaving digital traces does influence how people feel about those traces. Additionally, there is a tendency to trust organisations to use data within the law and individuals often gain benefits through incentives or because it serves their own self-interest.

Take the notion of the electronic butler, an intelligent software service which performs tasks according to stated or recorded preferences. A meeting is planned in the agenda for tomorrow in another neighbourhood. The butler is aware of a preference for travelling by train. It therefore checks the times of the trains, estimates how far the walk is from the nearest station and gives a short weather forecast. Once the individual has confirmed the details, the butler requests and purchases an electronic ticket for the journey. The following day, on approaching the meeting venue, the butler announces the arrival to the real or virtual reception and receives further guidance as to where the meeting will take place.

Preferences for tea, coffee and other services may also be communicated between the butler and the local system. Such a butler could be of great benefit to individuals, but it must store significant information and share parts of this with the other systems with which it communicates.

Improving healthcare records and monitoring of patients is another example of where there could be significant benefits to individuals. Instant access to a patient's full medical history when they first arrive in a hospital seems desirable. And it would be anticipated that such data would not be shared either with an employer or an insurance company.

5.5 How to become a consenting individual?

Will things really change from the current situation where we have grown used to the collection of personal data? Social and political institutions will not become less trustworthy, nor will the guardians of civil liberties become less vigilant. And experience of handling data privacy issues is growing all the time. So if *pervasive computing* were implemented, and with a minimum of legislative subtlety, one would assume that the core of existing data protection principles would be maintained:

- Individuals should be in control: they must give their consent and be aware that the data is being collected
- The purpose for which the data is collected and the limitations of its use must be "attached" to the data itself and this should be verifiable over time
- Appropriate anonymity should be maintained so that the user cannot be identified
- Systems must not link different user transactions without prior consent
- How long data is stored and considered relevant must be specified and agreed

Some technologists believe maintaining these elements is achievable, primarily through a notion of "sticky" privacy that binds data to specified uses within the critical infrastructure of *pervasive computing*, that is, at a very basic level of coding and naming procedures.

Others believe that the "on-the-record" environment presents a major challenge because the notion of personal choice or consent is no longer relevant. The exponential growth in the digital traces that individuals leave behind means it is no longer possible to control when or where information is collected or how it is correlated. As a minimum, there may be potential to allow consent at a system level, but this is again highly complicated. Travelling on public transport, for example, may imply video surveillance, transaction information and choice of route information being collected from named individuals. It may not be possible to opt in or out of this collection process, but one can be made aware that it is occurring and choose how to use the service accordingly.

What to do?

Three basic positions emerged as to how one could respond to the question of *pervasive computing* and privacy:

- *Pervasive computing* must be limited in its implementation due to the need to protect individual rights and privacy.

- It is possible to implement technology that enforces data regulation in this new environment: encrypt individual identities, limit access to information etc. The challenge is to awaken interest in this type of technology, which does not yet have a market.

- Acceptance that the current notion of privacy will be redundant and that a new way of dealing with privacy will replace it.

A single simplistic answer does not currently exist as to which of these generic approaches would be correct. Indeed, the expert group had a tendency to adopt a different approach according to the context under discussion. There was a consensus that there is much work to do to improve our understanding of what is worth fighting for within the current framework of privacy, and of how *pervasive computing* will influence this framework. It is also clear that the better the implementation of *pervasive computing*,

enabling clear standards and generating

clear benefits, the easier it will be to exploit technologies to protect existing privacy concepts and principles. Finally, building on these previous points, the better informed the debate on these issues, the more likely it is that misuse can be prevented and over-reaction avoided.

Antigone the ant meets a future object

I am really delighted to finally meet a real networked object. Amid all these words and wonderings, I may actually get the chance to understand first hand what is going on with this *pervasive computing*.

I am also delighted to meet you, but I fear that I won't be able to help as much as you want, especially with the idea that I can give you something first hand.

Oh dear that does trouble me, I have done my best to represent the ants in this little book, why can't you do the same for the objects?

Please do not misunderstand me, I will be delighted to represent the objects, my problem is that we are not the only thing involved. We get dragged into such a conversation exactly because we can be grabbed hold of, there is something irreducible about us. But the truth of the matter is that the reason why we are interesting is because we have become something more than we were before and that little bit extra is exactly what you cannot grab hold of.

Why not?

Because it is the bit of us that is connected up with a network, flying around in the ether, getting stored in computer servers and processed. It is the fact that we have become a new type of information carrier. But there is also something more than this. The human social environment has always been full of symbolic information systems that are embedded in our ways of going about things. Take money as an example. It is centuries since the weight of the precious metal in any way equated with the value of the coins an individual uses and now with electronic systems, physical money becomes increasingly irrelevant as payments are made with credit or debit cards, another type of object which again gets socially loaded with information according to its colour and from which bank it has been released. And of course there is also something special about humans and money as in some cultures collecting enough of it is equated with both happiness and the opportunity to look back and feel that one has led a good life.

Gosh, I had never thought of it that way. It is a lot more complicated than worrying where your next meal is going to come from.

Absolutely. The real point is that humans live in a world full of symbolic structures. They learn very early how to assign meanings to these and how to operate within them. Objects are never just objects, even when they are admiring nature for its naturalness it is already an informed social response.

Right. But have we not strayed from the

subject somewhat? I wanted to know how you networked objects might change things?

Perhaps you are right. I guess my point is that the networked object is the bit of the iceberg you get to see. What is also going to change is the nature of the social and symbolic reality within which humans live. Let's take the traffic light as an example. It is a physical object, already wired up, introduced to help regulate how cars and pedestrians get around town. But what happens when it becomes a more sophisticated information carrier? It can transmit information on cars that have gone through the red light and debit the fine directly from the owner's bank account. It can also count the cars passing through in each direction and thereby help regulate the next set of lights in a self-organising network. The traffic light may also share information with the car's navigation system, providing clues as to the most efficient route to work. So suddenly a dumb traffic light that goes red, amber, green in a set sequence, gets to decide when to change colour, penalise wrong-doing, collect the money and advise on the best route to work.

I see, it is as if objects can stand in for one another, doing each others' tasks because they are all connected. The thing-ness of the object is quite transitory.

That is how it seems to me. We will do a lot of regulating of behaviours in all sorts of ways. It is something the humans have not woken up to yet, namely that the protocols by which we networked objects operate are going to be really important. They offer mini-constitutions for how the structuring of social spaces will happen in the future. What used to go in the rule book will end up in the protocols governing how machines interact. And the exciting thing is that how these protocols play out will only be known when they start to interact.

Well, if the humans need reassurance that this can work rather well they need only look at my own species which creates a rather impressive level of order out of a lot of random activity. But I think I shall stick to the doing part of things. It is what ants do best. And I shall only start to worry if I pass humans hugging traffic lights the way some of their forefathers used to hug trees. Then we will know that their taste for symbolism has gone too far!

PVC Experts

Ulrich Brandenberger
Swiss Re
Risk Engineer
Switzerland

Douglas Dykeman
IBM Zurich Research Laboratory
Manager Services & Software
Switzerland

Matthias Haller
em. Professor for Risk Management
and Insurance Economics
Risk Dialogue Foundation
Switzerland

Stefan Hild
IBM Watson Research Laboratories
Senior Manager, S&A Solutions
Switzerland

Lorenz M. Hilty
Swiss Federal Laboratories for Materials
Testing and Research (EMPA)
Head of Technology and Society Laboratory
Switzerland

Walter Hirt
IBM Zurich Research Laboratory
Sensor Networks
Switzerland

Bruno Käslin
University of St.Gallen
Project Manager
Switzerland

Albert Kündig
Centre for Technology Assessment TA-SWISS
Member of the TA-SWISS Steering
Committee
Switzerland

Katrin Meier
Stiftung Risiko-Dialog
Leiterin Kommunikations-Technologie
Switzerland

Beat F. Schmid
University of St. Gallen
Managing Director Institute for Media
and Communications
Switzerland

Author
Lawrence Kenny

Core Team
Danielle Bütschi
TA- SWISS
Walter Grossenbacher
TA- SWISS
Walter Hehl
IBM Research Laboratory
Yvonne Neville
Swiss Re Centre for Global Dialogue
Walter Pletscher
IBM Research Laboratory
Karin Vey
IBM Research Laboratory
Markus Wichser
Swiss Re Centre for Global Dialogue
Alex Widmer
Swiss Re Centre for Global Dialogue

Publisher
Swiss Re Centre for Global Dialogue
Order: global_dialogue@swissre.com
www.swissre.com/globaldialogue