Same subject as GCA

Proposal for EPSRC Managed Programme in Ubiquitous Intelligent Systems

Version 6 26 May 2003

Comments to : Morris Sloman
Imperial College London
m.sloman@doc.ic.ac.uk

Thanks to the many people from EPSRC ICT SAT, UKCRC Ubiquitous Grand Challenges, UK-UbiNet and others who have provided input to this.

1 INTRODUCTION

Ubiquitous systems consist of large numbers of 'invisible' computers embedded into the environment which may interact with mobile users or form intelligent networks for sensing environmental conditions, animal or human activities in the natural and built environment [1]. Users will experience this world through a wide variety of devices, some they will wear (e.g. medical monitoring systems), some they will carry (e.g. personal communicators that integrate mobile phones and PDAs), and some that are embedded in buildings or in vehicles they use (e.g. car information systems). This heterogeneous collection of devices will interact with sensors and actuators embedded in our homes, offices, transportation systems to form an intelligent pervasive environment or 'ambient intelligence' which aids normal activities related to work, education, entertainment or healthcare. Sensor networks are a particular type of ubiquitous systems used to monitor the environment for security, to detect earthquakes, volcano eruptions, wildlife habits, weather conditions and military applications. In general, it is not feasible to provide wired communication to the large numbers of sensors, embedded or mobile devices so wireless communication is essential although access to wired backbone networks is also usually required.

One of the key attributes of ubiquitous systems is being aware of context which could include location and activity. The ability to track movement and activity raises serious fears over privacy and issues of what can be trusted in ubiquitous systems. However, the provision of context aware applications and services, which are able to determine and adapt to current activity, are key to supporting people in their daily activities. Seamless and ubiquitous connectivity with quality of service guarantees is needed over disparate communication technologies. Very low powered devices are required which can be powered from the environment e.g. using solar power or electrical pulses from the body when implanted into the human body. Autonomic, self-organising and self-managing systems are essential to cater for potentially billions of ubiquitous computers in everyday environment. Human factors arise to cater for new modes of interaction and 'invisible' ambient technology which must be useable by non-technical people. There is a need to support transient organisations and dynamic, potentially mobile work arrangements including virtual teams, ad-hoc collaborations and virtual organisations, but what are the sociological impacts of this ubiquitous, always connected technology? New business and accounting models are required for funding the intelligent infrastructure of adaptable networks, processing and storage services which must be available everywhere – home, work, shopping malls, in trains and planes.

There have been a few isolated projects in some of the above areas and considerable research in wireless technology, but there is a need to bring together electronics, wireless communications, computing, artificial intelligence, human factors and social aspects within a managed research programme to provide the impetus for interdisciplinary collaboration This convergence of different approaches needs to be supplemented by a consideration of a number of key theoretical questions. Ubiquitous systems challenges many of the principles that underpin the prevalent arrangement of users, computers and applications. If we are to realise ubiquitous intelligent systems in a principled manner then we need new theories to be developed in partnership with those who engineer this class of systems

There are many research challenges. One is to cater for *large* scale systems such as countrywide healthcare systems as well as to scale *down* for systems such as body area networks of sensors. Another is to ensure that the software structure, though highly interactive and mobile, remains under full control; errors, which are inevitable, must be tolerated and dealt with effectively at every structural level, and the systems must be able to report their activity. Such challenges entail new principles of engineering design, and new rigorous models for virtual machines that embody these principles. Such models are needed not only to assist design, but to support explanation. Pervasive ubiquitous systems entail a duty, greater than with present-day systems, to explain them both in scientific and in lay terms, both to our successors who will maintain and adapt them and to society which will increasingly depend upon them. These explanations involve new theories, indeed a hierarchy of models and calculi at many levels of abstraction, and tools for computer-aided analysis based upon these theories. Together the design principles, theories and tools will constitute a science of ubiquitous systems.

2 CONTEXT

There is considerable interest in ubiquitous systems with research programmes in Europe [2] [3] [4], as well as the rest of the world. However the UK involvement in European programmes has been comparatively low although we have considerable involvement in related areas such as wireless communication, networks, intelligent agents etc. A small number of UK institutions have played a key role in shaping the research agenda for EU initiatives such as the disappearing computer initiative, which focuses on the design of ubiquitous computing systems. It is clear that the UK has a leading European role in some elements of ubiquitous computing and that UK has a very strong track record in wireless and mobile communications which are an essential component for ubiquitous systems [5]. A coordinated UK programme would capitalise on the current UK lead in elements of ubiquitous computing by creating a research community with critical mass across a breadth of subjects.

The EPSRC Equator IRC [6] covers some of the issues of ubiquitous systems with a focus on understanding and support for new forms of interaction to emerge from ubiquitous systems. This programme would complement Equator by considering the broader middleware and theoretical issues in partnership with user-oriented concerns. Broader social concerns have also been addressed by past projects related to a Virtual Society funded by ESRC [7]. The DTI has recently started a new funding initiative – Next Wave Technologies which is meant to foster academic/industrial collaborative projects related to ubiquitous systems, but the focus is more on shorter term prototypes for industry [8]. An EPSRC sponsored academic network UK-UbiNet [8] has been established to raise awareness, coordinate research initiatives and to ensure that a healthy community of researchers emerges. Despites all the above activities there are still many research challenges to be overcome before ubiquitous systems become a reality. Moreover, these challenge require different research communities to work together. For example, the large number of distributed interactive devices means that builders of future communication infrastructures need to work in partnership with those who understand user interaction and those who can design ultra-low powered devices. A managed programme offers the opportunity to bring these currently disparate research communities together to address the broad set of challenges to emerge from ubiquitous computing. Two reports from the UKCRC Grand Challenges workshop focus on Ubiquitous *Computing* issues [10][11], but there are research challenges which cover other disciplines such as electronics, materials for sensors, human factors, social aspects and application areas relevant to healthcare, engineering, physics, biology and the environment. The applications of ubiquitous systems are relevant to all the other Research Councils.

A multi-disciplinary managed program is needed to bring together the researchers from computing, electronics, material, engineering, social science, psychology, business management as well as application domains such as healthcare, environment, biology and e-science. This will enable projects of the required scale to be mounted and will ensure that the concepts underlying ubiquitous systems become known outside the research community (e.g. in industry) in a timely fashion, enabling the UK to establish a lead in this area. It will also provide the leverage for additional funding from Europe and industrial sources.

Potential for international collaboration exists not only with the usual partners in USA, Canada and Europe but also with countries such as India and China which have a strong research base is wireless networks and are developing projects on context aware intelligent systems.

3 APPLICATIONS FOR UBIQUITOUS SYSTEMS

By their very nature, ubiquitous networked systems are likely to pervade all aspects of our life in the future. The following lists some typical applications. We would expect some projects to be focused on enabling technology but to include small application prototypes. However we envisage some larger application oriented demonstrator projects to also be funded. There are a few current projects in the UK or in other countries covering some of these applications (see [9] for links).

3.1 Healthcare

Ubiquitous systems provide the opportunity to dramatically enhance the quality of life of individuals with clinical conditions through healthcare provided in the community. The last decade has witnessed a rapid development of new monitoring and intervention medical devices including the monitoring of blood glucose level and controlling insulin delivery in patients with acute diabetes, multiprogrammable brain stimulators for epilepsy patients, cardiac monitors and defibrillators etc. Wireless networks in the home, buildings and cities would allow devices such as these and other sensors to interact with the centres of care to enable continuous monitoring and assessment of patients. Any discrepancies can be brought to the immediate attention of healthcare workers with access to specialist care if needed. This can reduce health crises, consequent hospital stays and help individuals with clinical conditions to lead a "normal" life. The ability to monitor, not just severe high-risk patients, but also the general lower-risk patients over long periods of normal activity, will facilitate research into trend analysis preceding cardiac problems, strokes or epilepsy. The availability of integrated patient information monitoring will therefore open up a whole new series of clinical research topics, where the fusion of objectively acquired multi-sensory data can lead to more reliable and earlier detection of life threatening events.

Another aspect of healthcare would be monitoring the activity of elderly people e.g. with dementia to determine when they need help. There is an need for systems which can monitor people to determine normal behaviour patterns or medical symptoms and then alert the patient or relevant authorities when something abnormal occurs. An intelligent system could monitor a person's usage of medication and warn if incorrect medicine was being taken or the person forgot to take the medication. All this is a key objective for Government to meet the needs of Care in the Community Agenda and reduce bed shortages in hospitals by means of early discharge. However deployment of ubiquitous systems technology for healthcare is partly conditioned by technical challenges (especially mobile monitoring) but possible more important, by organisational and culture barriers, so research into understanding these is critical

As well as healthcare, the area of fitness and well-being (not merely prevention of disease) is a potential application area. While running, swimming or training in a Gym, data can be gathered through a number of sensors, logged, and integrated into a regime for the individual. Research on this monitored information can be used to offer better advice (both to increase the effectiveness of a regime, and to reduce risk of injury).While there are fewer security problems than in healthcare, questions of privacy and integrity of data are still present.

3.2 Environmental and Wildlife Monitoring

Much of the original impetus for sensor network research came from military applications, as a consequence of the requirements of these application domains such sensor networks can be applied to monitoring tasks in many harsh environments. Sensor network features include self-configuring ad-hoc networking techniques which cooperate in filtering and forwarding information to base stations with wired network or more powerful wireless connectivity. Sensors may be distributed by hand, from vehicles or even dropped from planes. Intelligent wireless based sensor networks can be rapidly deployed in areas where needed to monitor water levels and flows for detecting floods, coastal erosion, seismic activity for earthquake detection or volcanic eruptions, as well as wildlife habits for research purposes (c.f. Great Duck Island Project in USA). Sensors based on mobile robots may be used in disaster recovery e.g. to detect trapped people in collapsed buildings or mines [12] which requires integration of communications, software engineering, vision capabilities, learning and the ability to understand possibilities for action in a structured environment.

3.3 Personal Assistants

Ubiquitous system technology can be used to produce intelligent personal assistant systems which can help users as they go about their normal activity. Such devices could combine the functions of a diary, remind the users of names of people they meet and possibly the situation in which they last met the person, be a set of secure keys, act as a credit cards. In addition they could hold medical and other personnel information which could be of assistance to the user/owner of the device or other professionals such as the doctors. A personal shopping assistant could suggest suitable clothes which will match the user's tastes and tone in with other items in her wardrobe or interact with intelligent food packaging to make sure the person does not buy items to which they are allergic or have dietary contradiction as well as suggesting alternative products within the user's budget. The personal assistant needs to be able to determine current activity, combining information from many diverse sources and provide 'intelligent' advice about travel planning, entertainment, restaurant menus to avoid food products which could cause allergic reactions etc. This application area requires collaboration between mobile computing, semantic web and cognitive systems communities.

3.4 Intelligent buildings

Intelligent building are already being wired up with sensors and actuators to control lighting, air conditioning, heating and elevators based on current occupancy levels. In the future, these sensors can interact with personal identification devices to open doors, detect intruders, detect items being moved by unauthorized people etc. Security is likely to be one of the key application areas for systems with sensors to detect motion, sound as well as opening of doors and windows. These sensors can be easily distributed throughout a building, and communicate with a central location by wireless links without the need for expensive and intrusive wiring. The central system can analyse and act on the monitored data and control locks. Novel speech or biometric identification techniques may be incorporated into these systems.

It is possible to conceive that in the future intelligent glass could be used in buildings to provide one-way mirror effects, when required, or window darkening so there is no need for curtains for privacy. The use of intelligent paint which could change colour would enable us to adapt rooms to mood or season e.g. cool colours in summer and warm colours in winter at the touch of a button.

The main issues for intelligent buildings and 'smart homes' are to do with systems integration and consumer acceptance of the new technology with respect to ease of use and cost-benefits.

3.5 Intelligent clothing

Wearable computers are being developed to permit access to computing services while walking around to support users in tasks such as vehicle maintenance, medical advice and recording for ambulance crews while remote from a standard computer system . This will be extended so that medical sensors could be embedded into clothing or jewelry to provide non-intrusive health monitoring. Including a simple passive radio activated (RFID) tag on clothing would enable you to query the supplier of a garment you particularly like the look of when you see someone wearing it or colour match items in the shop with items at home.

Medical sensors may be embedded into clothing or jewelry to provide non-intrusive health monitoring. Including a simple RFID tag on clothing would enable you to query the supplier of a garment you particularly like the look of when you see someone wearing it.

3.6 Intelligent packaging

Packaged food sold in supermarkets could have embedded intelligent sensors which monitor life of the product in the context of the conditions under which the food has been stored and transported to determine that it might not be suitable for consumption. The application of intelligent packaging has already started through simple tagging of parcels to track their passage through delivery systems. The main benefits are for better supply chain management but this could also benefit healthcare in terms of monitoring food being purchased and consumed to improve diet.

3.7 Traffic monitoring

Cities such as London are already using cameras for monitoring traffic for congestion charging and determining traffic congestion conditions. These are currently fixed cameras, at predetermined points, but cars could have cameras or other sensors fitted to determine local traffic conditions or pollutions levels then this information could be reported to coordination centers. These sensors could also be used to limit driving speed of a vehicle. This could be based on traffic, detect potential accident situations and provide warning or even apply brakes automatically.

Transport for people and goods can be greatly improved through pervasive information systems support. The extent to which it can automated (e.g. removing driver error in similar ways to those used by airline industry) is highly dependent on the availability of very reliable, and cheap control systems. To date, such systems are expensive, people intensive (c.f. air traffic control) and often intrusive. To make such systems affordable and invisible is a major challenge for ubiquitous systems.

3.8 Entertainment

The IT entertainment and gaming industry represent a significant worldwide market. Current estimates suggest that over 50% of Sony's current income comes from its investment within the computer gaming industry. The importance of this market is reflected by the major investment of Microsoft in the development of the X-Box. Similarly, phones are already being marketed as entertainment devices and third generation mobile video receivers are coming onto the market.

The importance of mobile gaming as a potential market place for future telecom operators has already been recognized by manufacturers such as Nokia who have established mobile gaming structures that span a range of media (www.nokiagame.com). The success of these initiatives have already stimulated a market for wide scale mobile games that combine a range of devices and exploit location and movement. UK researchers have played a significant role in developing initial explorations in this area. For example, as part of the Equator IRC collaborations with Blast theory have resulted in a mobile gaming system that has been nominated for a technical Bafta and has one the international Ars Electronic prize.

The emergence of new devices will allow us to support multiplayer games, interact with intelligent game agents within the network and to deliver entertainment content targeted to these devices. A coordinated programme will allow the development of new forms of content to take place in tandem with the construction of new forms of interactive devices.

3.9 E-Science

The growth in e-Science initiatives and the development of worldwide collaborative infrastructure for e-Science has led to considerations of how sensors and devices may be connected to a digital infrastructure. Ubiquitous systems offer the possibility of broadening the impact of the e-Science infrastructure to consider the development of future laboratories which sense the activities of the scientists who inhabit them and provide tracking and awareness of the activities that take place within these labs.

The development and deployment of widespread wireless networks allows us to consider the potential offered for those scientific communities who need to work in the field gathering materials or undertaking field trails of different forms. The challenge is how we realize these devices in such a manner that the inferences they make are accurate and dependable and to understand how best to annotate the scientific process with the recording of activity information.

4 RESEARCH ISSUES

Although there has been considerable progress in some areas such as the wireless technology to support mobile communications, there are many research issues which need to be addressed in order to realise a ubiquitous intelligent environment which is able to support people in their normal activities.

Electronics

Ubiquitous sensor systems present many challenges to microelectronic system and device designers to permit interaction with the environment through distributed micro-transducer systems. Fundamental among these is the need to develop digital and radio communication systems with ultra low power consumption; for example, batteries will be impractical for many applications, so energy-scavenging (extracting power from locally-available light, heat or vibration) will be the major power-source. A particular issue is the need for very low power devices and alternative power sources for devices. If you have 1000s (possibly 100k) of devices per person, you cannot keep changing batteries. Use of solar cells, fuel cells, heat converters, motion converters may all be possible. The biggest challenge is to design very low powered devices, transmitters etc although in the past our emphasis has been on faster chips. Development of new intelligent devices would bring together issues of MEMS, wireless communication, new sensors, low power electronics, and novel human interaction within low cost devices c.f. Motes research at Berkeley University [13]. Organic electronics could solve some of the human computer interaction problems for portable devices. New types of sensors are needed which can be used in very harsh environments such as volcanoes.

Systems powered by scavenged energy may have a mean power budget of around 100 microwatts. This will require radical new low-end microprocessor architectures that combine the power-efficiency of current 8-bit processors with the efficient programming models of current 32-bit processors, perhaps with enhanced signal processing capabilities for sensor signal conditioning. Radio communication will have to be rethought, as current sophisticated low-range systems (such as Bluetooth) require far too much power. Novel device techniques will be required to develop sensors and MEMS for energy-scavenging systems. All of these developments will have to be sympathetic to future process technologies where, although the small feature sizes offer very low dynamic power consumption, this advantage can easily be cancelled out by increased device leakage. Very low duty cycles, with efficient powering down of individual subsystems during idle periods (to minimise leakage losses), will affect processor architectures and radio protocols alike. New technology such as fully depleted silicon-on-insulator CMOS, combined with pass gate logic or vertical devices where the gate is 'all-around' a turret which contains a channel on each side could be investigated for ultra low powered devices. We also have to consider the trade off between hardware and software with respect to power consumption.

Some sensor systems will be used in harsh environments - monitoring volcanoes being a case in point - and this will place additional demands on devices and systems. Other applications may require sensors with propulsion and navigation systems, and again these will challenge designers for the foreseeable future.

Developing a UK programme for sensor electronic system design and device technology is very much in line with the major recommendations of the House of Lords report "Chips for Everything" [14] and would provide a major focus for the UK design community, paralleling the programme at UC Berkeley [13] that is developing "motes" (also known as "smart dust"). Clearly simply duplicating the UC Berkeley programme would be wasteful, and the UK programme should both establish collaborative links with the US work and develop its own unique direction.

Wireless Communications

Ubiquitous systems need to interact with mobile users or devices without wired infrastructure. This generally involves some form of wireless communication – Personal Area Network (PAN) technologies such as Bluetooth and ZigBee, and remote devices using a variety of heterogeneous networking options such as fixed line, cellular and WLAN technologies and, more recently, digital broadcast networks. Mobile users (e.g. for healthcare applications) may need to seamlessly switch between different modes of communication when they move from house to outdoors etc. New communication protocols are needed to limit power drain and provide reliable communication between unreliable devices in wireless sensor networks which are essentially ad-hoc in nature. The design of wireless communications systems has traditionally followed the principle of provision of continuous coverage at an acceptable level from a network operator, who is also the service provider. This has been the pattern for 2G services in which voice, and latterly text messaging, have been dominant. The situation changes dramatically if i) service providers are not necessarily network operators, ii) variable bandwidth services are offered and iii) universal mobility is not a requirement.

A wireless LAN is an example where coverage is no longer universal but can occur in islands which may or may not be connected. Therefore a particular session may not be continuous but is commenced or continued whenever a user is within range of a transmitter. This scenario has previously been envisaged in terms of islands of cordless or limited cellular coverage. It is necessary to extend this concept to a variety of wireless networks interconnecting a variety of intelligent user devices via intermediate edge, hotspot and core networks. Such networks and their component devices are created and extinguished in a dynamic way and offer a viable alternative to traditional fixed operator networks, for delivering enhanced electronic services for ubiquitous intelligent systems.

Mobile IP has been defined to support the ability of mobile devices to change their points of attachment without losing their session connectivity. This concept must be extended to accommodate wireless networks in which large numbers of individual devices may join and leave a network in a dynamic way. One of the most important issues is authentication, and the service provision of a scalable, unique, verifiable identity for users accessing dynamic local transport networks. It is key to security and trust relationships, to ensure, for example, that the user cannot deny the service access took place, and to protect the local transport provider in the case of malicious access attempts to services through their equipment. Local service providers also need to be protected against denial of service attacks to their network. On the other hand, users will not have a long term relationship with the local transport provider in the same way that they currently have with a network provider, and will need to use techniques to protect their information from the local transport provider. Some confidentiality techniques, such as IPSec, present difficulties when used over a wireless link since they deny access to some of the information used by wireless TCP techniques to improve throughput. Development of methods whereby trading and service discovery functions can be used to explicitly negotiate required parameters for secure channels rather than relying on side information are therefore an essential requirement in wireless communications systems.

Middleware

The traditional remote procedure calls or object invocations used in distributed and internet protocols are too heavyweight to be used for implementing ubiquitous systems which are essentially event based. Event-based asynchronous messaging may be more appropriate but distributed semantic cooperative filtering and expression of interest (subscription) paradigms must be supported by ubiquitous systems. This may require temporal and spatial locality tagging of events for some applications.

Self Management

New approaches are needed to provide flexible and adaptable software and hardware both for mobile devices and the intelligent environment. The scale of these ubiquitous systems necessitates 'autonomic' (self-organising, self-managing, self healing) systems which can dynamically update software to cater for new services and applications. Design tools and theories may be needed to support large-scale autonomic computing for small devices – current work is focused on web servers, databases etc. A related issue is discovery of resources/ services that your require. (see IEEE computer Jan 2003 pp41-50 for a good discussion on Autonomic computing). There is a need for self-stabilizing algorithms and techniques which will cope with large percentage component failures and even unsignalled failures.

Information Management and Provenance

Vast numbers of sensors can potentially generate petabytes of data which needs to filtered, aggregated, reduced etc. as close to source as possible e.g. using programmable networks, collaborative sensing strategies. Sensors may have to be programmed to look for specific events. How do you find required information – the information being published needs to be described using meta data, (not just time stamp, spatial location/origin) c.f. work on semantic web. In some cases it may be better to send queries to sensor networks rather than assume they send information to databases. Extracting the relevant information such as speech in noisy background can be very difficult. There is a also a need for information-theoretic approaches for obtaining and processing of information in networks of potentially unreliable sources, processing components and communication links. To achieve these goals we need to develop a new category of algorithms and data structures that have a *tendency* towards convergence and consistency. In some situations, data will be highly dynamic both in content and origin, with individual devices appearing and disappearing more or less at random – this provides a big challenge for self management, optimisation etc. Biological paradigms may also be appropriate.

In some ubiquitous environments, the user will be inundated with data. Much of this data will be filtered and deleted. Provenance could be an attribute for filtering. The authoritative source of information (i.e. 'pedigree') as well as how it has been modified or aggregated may be needed for some applications such as healthcare. There may also be a need to provide an audit trail of how information has been used (c.f. electronic fingerprints). This may be for legal reasons or related to trust and could require annotation of data with extra detail or comment. This will require design mechanisms which generate the meta-data annotations so that processing components can control where information goes, correctly report its origins despite any transformations, and prevent intermediate processing stages from compromising the meta-data.

Programming and Design Tools

New ubiquitous system programming models are needed to provide uniform naming and access to resources. The resource interface specification need more sophisticated specifications of behaviour than provided by existing distributed systems technologies. Systems must be programmed to work in the presence of failures without complex operating systems and communication systems to mask faults at lower layers. An aspect of the programming system challenge should be to find a minimal approach for the data structures and their semantics to allow decorative syntaxes which permit annotation of thread and communication with stochastic properties e.g. HOT Highly Optimized Tolerance theory which allows some interaction between intended and derived properties through introspection/reflection, as well as supporting the appropriate division and recombination of processing and data amongst individual under-resourced processing elements. Note that the reliability improvements achieved here apply equally well in the large, where a small-fold increase in redundant processing and storage would have massive pay-offs in availability, timeliness and, possibly, correctness. Other important targets for new programming systems are the improvement to design structures that yields intelligibility and (hence) maintenance and evolution and how to design systems of this scale for testability. Essentially there are two main problems to be solved – the techniques and tools to support the functionality of ubiquitous systems as well as ho to make them robust with respect to various kinds of failures.

Context Awareness

Ubiquitous systems cannot assist users without being able to accurately determine what they are doing. Sensor networks may need to be aware of what they are trying to measure or detect – human activity, cars, specific animals etc. Context awareness could include current geographical/spatial location, direction of motion, nearby neighbours, device capabilities relating to I/O, power, processing, activity of users etc. and it is a key property for many applications of ubiquitous systems. The observations obtained from sensors are subject to errors and ambiguities and require careful interpretation and assimilation to physically valid models in order to provide useful input to applications. Context modelling will draw on experience from robotics, computer vision and several areas of physical science but there are issue of how to represent and store context to be addressed. Research on context modelling systems and on the sensor systems should be linked to ensure that the sensors provide adequate valid data. A key challenge is obtaining the information needed to function in a context-aware manner. In some cases, the desired information might already be part of a user's personal computing space. For example, that space might include schedules, personal calendars, address books, contact lists, and to-do lists. Systems must sense more dynamic information in real time from the user's environment – such as position, orientation, people's identities, locally observable objects and actions, and emotional and physiological states.

Trust, Security and Privacy

Ubiquitous applications may involve collaborations between ad-hoc groups of devices, agents or infrastructure servers. It may require migration (mobile agents) or downloading of code. There are complex issues associated with knowing what resources, services or agents to trust. For example, do you trust your neighbour to send you packets for ad-hoc routing? This could be a denial of service attack to deplete your battery. From trust, and possibly recommendations, you may be able to derive authorisation policy as to what access you will permit to your resources or possibly what services you should refrain from using. There is a need to validate credentials without access to network infrastructure and certification authorities. Most ubiquitous systems have the capability of tracking users and determining patterns of activity – this can be very dangerous if it falls into the wrong hands. There are many problems to be solved relating to the trade-off between the desire for anonymity in some situations and the need for identity based tracking of context or authentication for commercial reasons such as credit card payment. We need logics and languages to be able to reason about these concepts and safe languages with robust typing amenable to analysis. Biologically inspired concepts such as artificial immune systems may be relevant. The research needed is not just technology solutions but also sociological aspects of trust formation when interacting with devices or people with a ubiquitous environment.

Location and context monitoring is already widespread e.g. video cameras surveillance, mobile phone tracking. New approaches to privacy are needed, such as user based controls on access to stored information, which can then be used to prove that you are not or were not at a specific location without revealing where you are. Overrides of controls for national security, legal or emergency services use are needed.

Ubiquitous Systems Theory

Ubiquitous systems will be large-scale, very complex, dynamically self-modifying systems. Current implementations are based on opportunistic ad-hoc implementation which rely on the skill of the programmers to make them work We need to develop a coherent and rigorous informatic science to provide the concepts, calculi, theory and automated tools to allow descriptive and predictive analysis of these systems at each level of abstraction. Two factors create a need for robust theoretical understanding beyond that required for most existing software systems.

The first is the larger range of relevant concepts; among them are locality and mobility, security and privacy, boundaries (physical and virtual), resources and trust, belief and knowledge, nondeterminism, stochastic behaviour, self-organisation, strategy, evolution, reflectivity. No single model can accommodate these; we need to develop a hierarchy of models, each with its own theory, and theories of how each model is realised or implemented in models lower in the hierarchy. For example, one model may embody a logic of trust or a calculus of planning, and this is implemented in terms of action disciplines in a lower model. The lowest models in the hierarchy should be the various programming languages used in defining subsystems, as well as the protocols governing all communications.

The second factor is the fluid, evolving, self-modifying quality of ubiquitous systems, both hardware and software. The failure to manage this feature of classical software systems, in terms of either concepts or engineering principles, has led to mountains of inscrutable legacy software. For ubiquitous systems the legacy will be worse: systems will become inoperable; and the damage will be worse because the systems pervade our lives more deeply. To defend against this, existing theories of mobile systems must be greatly enriched, supported not only by the existing discrete models of computer science but also by hybrid and stochastic models, together with scientifically informed automated software for predictive analysis, model-checking and simulation.

It is necessary to model *boundaries* (whether physical or virtual) and *movement* of components in order to understand ubiquitous systems. Calculi and logics are emerging for modelling, boundary access, process and boundary mobility, and message and host failure. Further research is needed on developing calculi and logics, together with associated programming languages, verification and performance tools, and assessing them within practical scenarios in order to reveal further challenges to the fundamental theory. Mobile users and the mobile agents supporting them may need to acquire local resources from the environments through which they move. Methods based upon logics and types now exist to *describe* allocation and deallocation of resources, such as storage and processing, in a dynamically changing environment where nodes can join or leave the system at unpredictable times; but methods to *control* resource usage and to *analyse* resource descriptions are in their infancy. There is significant potential for collaboration between theorists and practitioners associated with querying semi-structured pervasive information and discovering services using pattern matching and mobile query languages. It is important to establish trust via the verification of security protocols, and the security and privacy properties of software used for ubiquitous systems. Model-checkers and symbolic techniques based on logical proof, bisimulation or game semantics can be used to verify security protocols. In future, we expect the verification of implemented protocols, not only their abstract descriptions. Safe languages are needed with robust typing and static analysis to protect against buffer overruns. In order to support safe mobile code, programs should be distributed with proofs of their type safety, which can be checked once they reach their target destination. The next step is to provide formal policies that also guarantee resource-bound and access control properties of certified code, plus certifying compilers than can propagate such source properties and evidence for their validity to low-level code.

Human factors and Social Issues

While there exists considerable human factors guidance for designing various kinds of interface (e.g., graphical displays, direct manipulation interfaces, multimedia systems and Web-sites), this has largely been aimed at displaying information on desktop computers. It is not certain that these current approaches are appropriate to the design of interfaces for ubiquitous systems which extend beyond placing a monitor, keyboard and mouse on a desk. In particular, it is not clear that they are appropriate when computing functionality becomes embedded into the surrounding environment, into specific physical objects, or into objects that are worn.

Ubiquitous systems will need new models of human-computer interaction that are appropriate to embedding large numbers of interactive devices into the environments we inhabit. These models need to be sensitive from the outset to the fundamentally social nature of the environments that ubiquitous devices are placed and the subtleties involved in the situated nature of this technology.

New design methods will be needed that move beyond HCI techniques that have been adapted from sociology and psychology, and include new techniques such as those used in art and design traditions. The long term goal will be to establish an armoury of design methods to support the design of new technologies that are usable, stimulating, engaging and account for actual social practice.

New evaluation methods will be required to understand and assess future ubiquitous systems These will need to blend experimental techniques adapted from psychology to understand detailed device interactions and techniques adapted from sociology that observe the use of new technologies in everyday environments. As ubiquitous systems reach beyond the office and work these approaches will need to be supplemented by techniques from broader design traditions that gauge people's reactions to and feelings about new technologies

This methodological work will need to be complemented by broader considerations about the overall socio-technical arrangement that ubiquitous systems will be imbedded in and the implications that emerge for the management and governance of this class of system. What are the future policy issues to emerge from the widespread deployment of ubiquitous computing and the consequential ethical implications of these systems?

Business Models

Ubiquitous applications often assume a support infrastructure which provides storage, processing, network connectivity and various services in the home, office, on planes, trains and even in the street. How is this paid for? What is the business model which will make sure this expensive infrastructure is available everywhere.

Research is needed on how ubiquitous systems can be integrated into the business process and what changes are needed to support clients interacting with commercial ubiquitous intelligent environments.

5 Involvement of Other Councils

Most of the other councils could be involved from an application view. It is not obvious to what extent fundamental research input is needed from the other councils, although social issues are within the remit of ESRC. This research activity requires engagement of computing, human factors, communications, electronics, sociology, business management and possibly materials communities. If application oriented projects are included, NERC and MRC might participate

6 Scale of Activity

About £10M funding over 5-6 years, including input from other councils, assuming 3 calls.

The scale will depend on how much buy in we get from programmes other than EPSRC ICT.

We envisage research in Ubiquitous Systems continuing beyond the life of this programme with funding from normal responsive mode sources, EU and industry.

7 References

[1] XEROX / Mark Weiser "Ubiquitous Computing" site http://www.ubiq.com/hypertext/weiser/UbiHome.html

[2] ISTAG Report on Scenarios for Ambient Intelligence in 2010 ftp://ftp.cordis.lu/pub/ist/docs/istagscenarios2010.pdf

[3] ERCIM News on Ambient Computing http//www.ercim.org/publication/Ercim_News/enw47/EN47.pdf

[4] EU funded Disappearing Computer Initiative http//www.disappearing-computer.net

[5] Mobile VCE http://www.mobilevce.com/index2.htm

[6] Equator IRC http://www.equator.ac.uk/

[7] ESRC Virtual Society Program http://virtualsociety.sbs.ox.ac.uk/intro.htm

[8] DTI Nextwave technologies and markets http://www.nextwave.org.uk/index.htm

[9] UK-UbiNet: A UK Network for Ubiquitous Computing
http://www-dse.doc.ic.ac.uk/Projects/UbiNet/

[10] UKCRC Grand Challenge Proposal 2: Science for Global Ubiquitous Computing
http://umbriel.dcs.gla.ac.uk/NeSC/general/esi/events/Grand_Challenges/proposals/Ubiq.pdf

[11] UKCRC Grand Challenge Proposal 4: Scalable Ubiquitous Computing Systems
http://umbriel.dcs.gla.ac.uk/NeSC/general/esi/events/Grand_Challenges/proposals/US.pdf

[12] RoboCupRescue http://robomec.cs.kobe-u.ac.jp/robocup-rescue/

[13] Berkeley Smart Dust Project http://robotics.eecs.berkeley.edu/~pister/SmartDust/

[14] Chips for Everything: Britain's opportunities in a key global market, House of Lords Science and Technology Select Committee report, 28 November 2002.
http://www.publications.parliament.uk/pa/ld200203/ldselect/ldsctech/13/1301.htm