UK COMPUTING RESEARCH COMMITTEE

Ubiquitous Computing

Comments to the POST

UKCRC Response to POST Enquiry on Ubiquitous Computing

The UK Computing Research Committee (UKCRC), an Expert Panel of the British Computer Society, the Institution of Electrical Engineers and the Council of Professors and Heads of Computing, was formed in November 2000 as a policy committee for computing research in the UK. Its members are leading computing researchers from UK academia and industry.

Ubiquitous Computing Systems (UCSs) will consist of millions of potentially inter-linked computers embedded in clothing, people, animals, vehicles, buildings, in the countryside and in the atmosphere. It includes an infrastructure of networking and major computational facilities supporting wired and wireless communications, processing and storage. People will also be a part of UCSs, some concerned with their operation, others just as users. The UCS could, for example, help them to perform their normal daily activities, monitor their medical problems, entertain them and support collaboration and interaction between them.

The characteristics of a ubiquitous computing system are potentially billions of devices and sensors producing vast quantities of information. Many elements are mobile and communicate via wireless links. A UCS is generally considered context-aware in that it can adapt to current context – a concept which covers more than just location, since it can include available resources such as battery power, device capabilities, wireless bandwidth, the presence and activity of both humans and other entities within the UCS. The UCS will be very dynamic as entities are both mobile and adaptive but its introduction is likely to modify the behaviour of the humans which will use and interact with it, causing it to further adapt to the human activity. It can thus be considered a form of ecology which evolves to support people.

The potential applications include healthcare by monitoring chronically ill and post-operative patients in hospitals and facilitating early release to be monitored in homes. Closed loop control using therapeutic actuators such as cardiac defibrillator, nerve stimulation, drugdelivery will be feasible. Monitoring in the home will not be confined to chronically ill but for healthy people who are at risk due to possibly inherited gene defects. In addition elderly or disabled people could benefit both from monitoring and a UCS which assists them go about normal activity. Other applications include environmental monitoring for climate change, bad weather warnings, earthquake warnings etc. A sensor network monitoring water flows and levels along rivers could provide early warnings about floods. Interactive entertainment, commerce and education will become available everywhere and at any time.

Identifying potential applications is not an issue but there are considerable problems with actually implementing UCSs to make them a reality in the future. Engineering problems relate to scale, mobility, infrastructure, dependability, security, and power; there is very little understanding on how to build adaptive self-organising systems; there is insufficient theory underpinning the design and analysis of such systems; it is not clear how people will interact with the UCS; and there are concerns about privacy when systems are capable of monitoring all your activities.

This document highlights the issues relating to both Engineering and Theory and suggests that solving these problems will be a grand challenge for Computing research for many years to come.

Engineering Issues

Billions of entities – devices, sensors, mobile phones, software components, databases, people etc will be part of future UCSs. They will form dynamic communities for specific purposes. Many will be mobile and error-prone and some will be malicious and try to harm other entities around them. Such a fluid, dynamic system has to be autonomic – self managing, self-configuring, self-healing, self-protecting – but must cater for human intervention where appropriate to set preferences or override autonomic decisions where appropriate. Yet the subject of autonomic systems is in its infancy, as yet far from capable of dealing with systems of the envisaged complexity, or with the range of threats to their dependability and security that are likely to arise in practice.

Determining context in terms of location and human activity is still not a solved problem. There is no low-cost location technology that will provide precise, reliable location determination both indoors and outdoors. Determining human activity such as walking, sitting, driving, cycling, in a meeting, in the cinema etc. is difficult to do accurately. Will the system be able to ‘learn’ what people are trying to do or determine unusual non-foreseen situations and adapt accordingly? The technology for this is still a very long way off.

Provision of suitable power sources for wireless sensors in or on the body, in the environment or even in the home is a major problem. In many situations it is not practical to regularly change batteries – particularly if there are potentially thousands of sensors per person. Even in homes and offices, it is not always practical to provide access to electric power for large numbers of sensors. Use of power sources based on solar energy, motion converters, fuel cells, heat converters etc still requires considerable research and development. Low power techniques are also needed for processing and communication in order to conserve the limited available power, for example, the current internet communication protocols are not suitable for interactions among millions of miniature wireless nodes.

The vision of miniaturised communicating nodes which can easily be sprinkled about the environment or implanted in people like ‘smart dust’, has not been realised yet and is probably many years off. Work is needed on low-powered, communicating devices which can be unobtrusively embedded in people, animals, clothing and buildings. However, very simple embedded devices will need support from an infrastructure to provide communications, more complex processing and storage. Who will provide and manage this and how will it be paid for?

Ubiquitous systems will generate vast quantities of information from monitoring the environment, people or vehicles and as people interact with each other and services using these systems. Information will be have to aggregated, filtered and possibly anonymised. The information delivery system has to cater both for queries on information and delivering notifications of events or warnings (e.g., your blood sugar level is low). Making sure only relevant information is received by entities is increasingly becoming a problem in the internet and this will be compounded with mobile ubiquitous communication. We need to be able to determine the provenance of information, from possibly multiple sources, and how it has been modified due to filtering, aggregation or any form of intermediate processing.

Ubiquitous applications will involve interactions between ad-hoc collaborations of mobile devices and people as well as downloading of code to devices. How can you determine who or what to trust in such an ever-changing environment? Based upon predefined trust, recommendations, risk evaluation and analysis of past interactions, an entity may derive new trust metrics and authorisation policies for what access it will permit to its resources, what services it should refrain from using, or what security mechanisms (such as encryption) to use. The system is capable of tracking human location and activity at all times. It may be logging sensitive medical information. These privacy concerns raise both technical and social issues of how to control access to information. What information should be made available about a person may depend on their current context.

People will need to interact with the ubiquitous system or use it to interact with other people as they go about their normal activities. This implies that many different modes of interaction and languages need to be catered for, such as voice and gesture for input, with voice, visual and motion (cf. gesture) used for outputs.

The overall complexity of such systems are such that one has to assume that their software, hardware and communications will suffer from faults – these may be accidental, or the result of deliberate attempts to subvert or damage the system. Self-healing capabilities are needed to mask or recover from the effects of these faults, particularly as such systems become more integrated into daily life so will people’s dependence on them increase. Although there are techniques for developing secure and dependable systems, most large-scale computing application have proved to be notoriously unreliable and insecure. There are also problems of being able to identify all the requirements for such complex systems or foresee all situations and circumstances in which they will be used. Making systems which can adapt to many unforeseen requirements and situations is beyond our current engineering capabilities.

There is thus considerable research and education needed to allow the engineering of ubiquitous computing systems as envisaged above. Engineers building bridges or 100-storey sky scrapers have a well understood theory underpinning analysis of loads and forces, even under extreme conditions such as earthquakes. Engineers building experimental UCS rely on ad-hoc ‘try it and test methods’. There is very little theory to underpin tools and techniques for UCSs.

Theoretical Issues

Understanding, analysing and predicting the dynamic adaptation of heterogeneous systems to context and interaction with humans; and how humans react to adaptation within the system is beyond our current capabilities. We need to perform analyses to gain knowledge and assurances on such unknowns. We must develop foundational methods for design and engineering of UCSs based on clear and objective models of correctness, utility and reliability. Such issues are notoriously hard to define and design for, and become harder from the perspective of different individuals in varying contexts, and with varied experiences and expectations.

One of the first theoretical challenges is clearly to develop a rich conceptual framework for the ways in which entities interact and move among each other. Formal behavioural models and logics are now emerging in which the physical space occupied by systems is treated as one with the virtual spaces and structures of software and data. The theory has yet to encompass the interaction between potentially mobile cooperating self-managed entities, each of which adapts its own configuration and behaviour to cater for changes in its context.

In order to facilitate the development of practical applications, models are needed that capture both continuous and discrete dynamics (e.g., time and location). Such models, called hybrid, have received some attention, but they are still rather limited. For example, the best known existing tools allow time as a continuous variable, but not space, making it impossible to model applications such as, say, a driverless traffic system, where both quantities vary as a continuum. A concrete high-impact objective here is to develop tools and techniques to handle the continuous dynamics of hybrid models described by differential equations.

Stochastic elements arise in many aspects of complex dynamic systems of autonomous components, such as ubiquitous computing. Examples include the probability distribution for temperature at a particular location, or for an entity to be at that location at a particular time. Behavioural, spatial and hybrid models for ubiquitous computing also need to cater for such aspects.

Models are being proposed for the analysis and understanding of paramount issues related to trust in the ubiquitous scenario, where identity is difficult to establish and context can change very swiftly. Logics have been proposed for expressing trust in terms of notions such as belief, authority and reputation, and techniques have been studied to enforce and monitor trust policies. This gives the basis for a discipline for the propagation of trust, yet to be formalised in a satisfactory way or implemented in practical applications such as ecommerce.

More work is needed on developing suitable mathematical models from algebra, logic, statistics, the differential calculus, etc., to provide the basis for software tools. Such tools are indispensable for the design and analysis of complex systems, and have in fact been the object of intensive investigation for current programming languages and paradigms. Any languages and paradigms proposed for UCSs should be based on such models which will need validation through experimental research. It is unlikely that any single framework will cater for the complexity of UCSs, and this in turn calls for the ability to make use of multiple models in an integrated fashion. This is not easy to comprehend or do. For example, a system may be described in a behavioural model, and a desired property of the system expressed in a logic of trust – we will then have to make these work together.

Models per se are not much use without the tools to support engineers in their tasks of designing and analysing UCSs. As we mentioned, languages are needed that can be used to program and build these systems, to express and then to prove that they behave in the ways we expect without the instability that can easily arise in very dynamic self-adapting systems.

However it is not just engineers and computer specialists who will ‘program’ the UCS. Nontechnical users will also need to configure them and modify their behaviour to adapt to their requirements. This type of programming may also have to be in the form of spoken language and gesture in some circumstances. They will also want to be able to understand, predict and analyse behaviour without using formal theories and tools.

The Ubiquitous Computing Grand Challenge

Ubiquitous systems have the potential to revolutionise the lives of chronically ill, elderly, or disabled people and greatly assist the rest of the population in their everyday activities of learning, work, entertainment and travel. However, there are considerable challenges relating to engineering, the underlying theory and the sociological issues which have to be solved before the hype becomes reality. Many of the problems to be solved are rich, deep and in many cases the problem itself may not be fully understood. One approach is for multidisciplinary projects to bring together the engineers, scientists, psychologists and sociologists to work together on building and testing partial solutions relating to healthcare, transport, education etc. The unexpected inadequacies of these prototypes will help us understand hidden complexity in the problems and from these to develop the required techniques and tools for building future large-scale ubiquitous computing.

For the above reasons, UKCRC has identified Ubiquitous Computing Systems as a Grand Challenge (see http://www-dse.doc.ic.ac.uk/Projects/UbiNet/GC/index.html) whose goals are:

- To define a set of design principles that pertain to all aspects of ubiquitous computing; are agreed among both academic and professional engineers; are taught regularly in Master’s Degree courses; and are instantiated in the design and rigorous documentation of several computational systems with a successful operational history.
- To develop a coherent informatic science whose concepts, calculi, models, theories and tools allow descriptive, explanatory and predictive analysis of ubiquitous computing at many levels of abstraction; to employ these analyses to derive all its systems and software, including languages; and to justify all its constructions by these analytic tools.
- To develop ubiquitous computing methods and techniques that are sensitive both to the needs of individuals and society and the impact upon them. This will include new forms of interaction and new interaction paradigms for the realisation of human experiences that make ubiquitous computing useable by all.

It is only by combining engineering with science, while still addressing the social aspects in exploratory projects, that we will make progress in the Grand Challenge of Ubiquitous Computing Systems.

UKCRC would be happy to answer follow-up questions on any of these points. It would also welcome further more detailed consultation.

Morris Sloman
m.sloman@imperial.ac.uk
for UKCRC, December 2005