

# Sentient Computing?

Andy Hopper  
*University of Cambridge*

Sentient computing is the proposition that applications can be made more responsive and useful by observing and reacting to the physical world. It is particularly attractive in a world of mobile users and ubiquitous computers.

---

## 1 Location Sensing

Cheap sensors make it possible for computer systems to react to the physical environment. Sensors giving location information are probably the easiest to construct and deploy. Use of such location information makes it possible for user interfaces to be based on space itself. Such context-aware, or sentient, interfaces and applications have been constructed and used for a number of years.

Sensors tell us about the location or position of things. To reflect the requirements of different applications, we take three different approaches to categorising the concept of location. First, containment is where we say that an object is within this container, e.g. a room. Second, proximity is where we register that we are close to something. Finally, co-ordinate systems provide a point location in space, subject to some error value. These categories are not hard and fast and can blend together. Small containers are very similar to a co-ordinate system, and proximity has much in common with the concept of containment.

### Infra-Red Location

15 meter range  
diffuse  
room-scale accuracy 95% of time  
containment location



Fig. 1 Containment: Active Badge

Our first experience of developing a sensor specifically to provide spatial information originated in the late 1980s in the form of the Active Badge (Fig. 1). Personnel and equipment could be tagged using the Badge, which transmitted a unique infrared signal every few seconds. The transmissions were diffuse and receivers in a room picked up the signal, giving room-scale containment. It told us who and what was in which room. The Active Badge was the inspiration that started us on this whole line of enquiry.

In the case of proximity, promising commercial systems are starting to appear. The radio-based Bluetooth system gives accuracy of about 10 meters using the received signal strength indication (RSSI). This will improve to about 50 centimetres in future implementations by using specialised on-board ranging circuitry. Similarly RSSI information from Wavelan (802.11) systems together with heuristics about movement of people can be used to provide in-building location information.

Outside, the Global Positioning System (GPS) can be used which has given rise to a large number of applications. GPS is accurate to around 30 meters most of the time, although greater precision can be achieved, and is one example of a co-ordinate based system.

In order to test the impact of fine-grain location information we have developed a co-ordinate system for indoors. This uses a tag, which incorporates ultrasonic transmitters and an array of ceiling-mounted detectors. A detector on the far side of the room will register a pulse later than a detector directly above an object. Using this differential timing information, we can calculate the position of objects to within a few centimetres almost all the time (Fig. 2). If two transmitters are attached to a rigid object it is possible to compute its orientation. The Active Bat technology is likely to remain the basis of the most precise indoor location systems for the foreseeable future. There will be many applications that do not require this level of precision and refinement. However, as a research tool, it is providing us with valuable information on what can be done with very precise positional data.

The Active Bat system requires a substantial amount of infrastructure, particularly

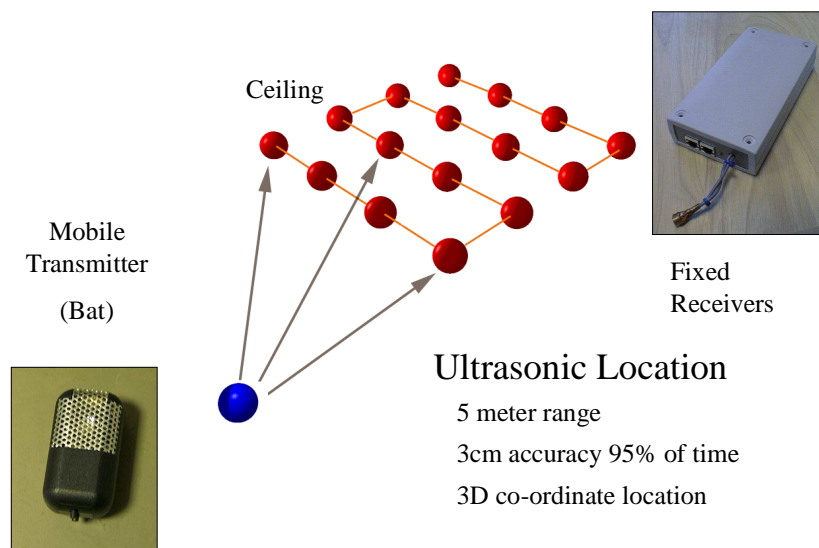


Fig. 2 Co-ordinate: Active Bat

in ceilings. A new technology, which may provide similar location information, is ultrawideband radio. By emitting very short pulses of several picoseconds duration we can measure propagation delays accurately at the receiver from transmitters spaced up to 20 meters apart. A large spectrum is used, for example from 3 GHz to 10 GHz, but the power levels are such that interference to other users is minimised. Ultrawideband transmissions may be less susceptible to interference in particular parts of the band and thus instrumenting buildings may prove much easier than with the Active Bat system. However it is likely that the precision will be some 10 times worse than the ultrasonic Active Bat with a location accuracy of about 30 centimetres most of the time. It also remains to be seen what the local effect of monitors and other metallic objects is on precision.

## 2 Spatial monitoring

Our sensors provide raw spatial facts about objects. They tell us where an object is, and possibly the direction in which an object is pointing. Location-aware applications need more than raw spatial data, they need to be notified of spatial relationships between objects that are significant for the execution of the application. But how do we decide whether a spatial relationship is significant? The approach we have adopted operates on the basis of zones of containment surrounding objects. In Figure 3(a) X represents a person and K a keyboard. Now suppose we have an application that needs to be notified when person X is in a position to use keyboard K – when X is possibly “holding” K. If the zone of confinement of K overlaps the zone of confinement of X, then the holding condition is held to be true and the application receives the appropriate space location event. The situation in Figure 3(b) indicates how this principle could be applied to support a multi-camera video conferencing system, giving participants the freedom to look in different directions while talking, or even walking around their offices.

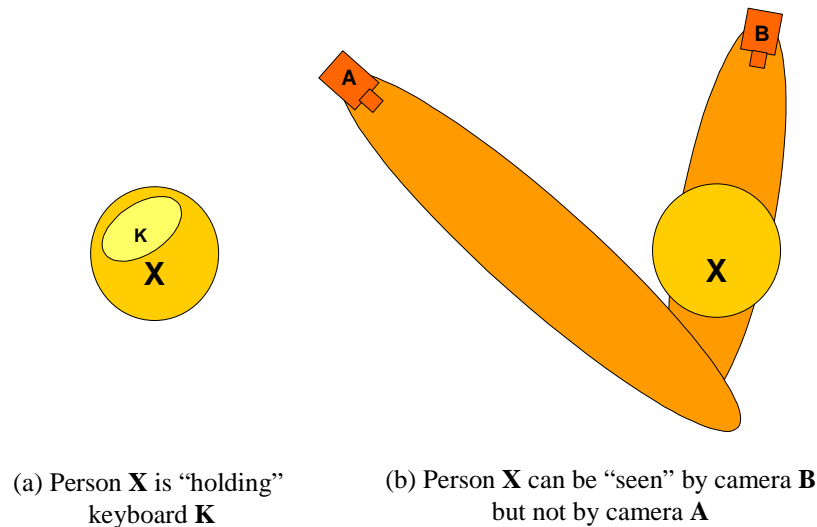


Fig. 3 Evaluating Spatial Facts

The principle of turning raw spatial data into application-significant events through geometric containment and overlapping is reasonably straightforward. Scalability can be addressed by applications indicating the interest and precision required. The computations are then only performed to the required level and the computational task scales linearly with the number of overlapping spaces. This approach can be thought of as the mouse/desktop metaphor mapped onto the physical world in real time.

The operational system that has been built uses a variety of sensors; allows space representations to change quickly; provides an appropriate governing event logic; uses caches and proxies to handle large volumes of data quickly; and executes in real time to satisfy a human in the loop.

Note that Figure 3 is a 2D representation of what in reality would be a 3D environment. This simplification can be made because, in general, people and objects tend to remain relatively fixed in the vertical plane. At the heart of such spatial monitoring systems we need to define a world model, which is easily understood by the user yet computable by the system. Is 3D important or is 2D satisfactory for most office and home applications? What is the precision of location information required? How can the spatial metaphor be made obvious to the user?

### 3 Data distribution

Publishing sensor data, which relates to the position of people and objects is one end-application. Beyond this we consider the automatic control of the digital environment with reactive and possibly predictive features. An attractive application for a user in a networked environment is the ability for the personal desktop to follow the user to any nearby device. In order to achieve this in addition to location information we need a platform for connecting and displaying information on all these devices in a ubiquitous way.

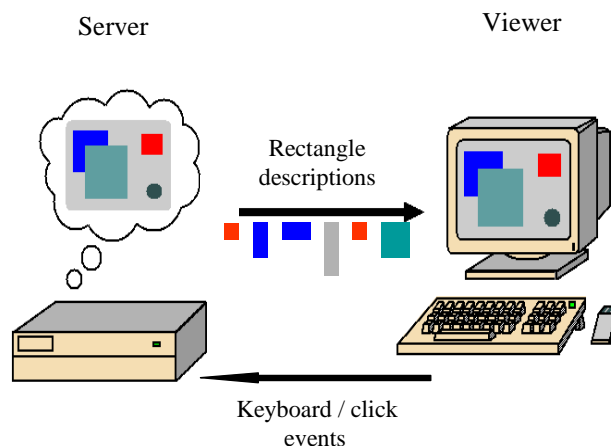


Fig. 4 VNC - The Platform

One way to do this is to tunnel connections to all devices using a simple device-independent protocol. We have devised one such ubiquitous platform called the Virtual Network Computer (VNC). In our approach the viewer, at the receiving end of the connection, has no state, and simply displays information graphically. The connection from viewer to server is also stateless, just keystrokes and pointer clicks. Our viewer is a particularly simple version of the so-called thin client (Fig. 4), with all application state and processing centralised on a server.

The absence of application state at the viewer eliminates any requirement for re-synchronisation, and the appearance is of user-interface mobility. In order to achieve this we have traded bandwidth, or more precisely we have relied on ubiquitous connectivity and low latency end-to-end.

The low-level nature of the protocol is the key to device independence, providing a platform that supports the connection of any device to anything. The connections can be one-to-one (fixed or mobile), and the streams can be split giving one-to-many, many-to-one, and many-to-many.

The performance of the VNC system has turned out much better than expected. By using a variety of compression schemes and caching it has been possible to operate useably across links with capacities of only 10's of Kbps and latencies of up to 40 milliseconds. Therefore incorporating the simplest devices with wireless connectivity within this framework now appears plausible.

## 4 Applications

Location information appears to be a powerful tool in constructing new applications. Opening and closing doors automatically is an obvious example. In our ten or more years of being immersed in such systems some of the most enduring applications have been those where raw location data is processed in a simple way and made available ubiquitously. A textual indication of where someone is, how fast they are moving, how long they have been there, has proved the most popular. Showing the local context, including who and what else is nearby, is also attractive. Publishing such information to the local (trusted) peer group saves time; if someone is not observed by the location system they are not available, whatever the reason. Graphical representations and in particular maps appear attractive but can become cumbersome in what is a familiar physical environment. So simple sensing and simple logic appear to work and applications presenting these stand the test of time. The containment location information provided by the Active Badge is quite sufficient for this purpose.

Personalisation by teleporting VNC desktops has also proved popular. The teleport can be triggered without using location data, but having a personal tag with a button, which acts as a personal ubiquitous controller, is neat. More precise coordinate location information as provided by the Active Bat becomes important for its ability to select the correct workstation or other device. Another use that takes advantage of the more precise location information associates a control function with any 3 centimetres cube of space. Typically this is done on the surface of a wall or other planar object and is normally a control trigger of some type. This use appears to have some merit and the walls of our laboratory are sprouting a number of such "active posters". For this specific

application a local proximity RFID tag can provide the same location information in a simpler way.

Specialist applications, for example surveillance where the selection of a particular camera is based on spatial data, can provide opportunities. However there is always scope for such bespoke solutions.

## **5 Observations**

*So what are the results of over a decade of research and what is the prognosis for the future? The research area is now very popular and is variously labelled ubiquitous, pervasive, ambient, calm, as well as sentient. In this paper we have permitted ourselves to give users attributes such as “holding”, “seen”, and even suggested the notion of prediction. Is this realistic?*

*We have learnt many aspects of how to construct such systems. Sensor information can be generated on a reasonable scale and presented to users in various ways. It seems the more direct the presentation the more attractive (or perhaps less irritating) the application. Some simple logic to interpret the data can be useful. Occasionally a domain-specific agent operates as envisaged.*

*Our attempts at automatic control without user intervention have not proved enduring. For example automatically teleporting to the nearest screen throughout the laboratory did not stand the test of time. Similarly automatic routing of phone calls had sufficiently serious flaws that the human operator remained as the interpreter of location data. User service profiles were attempted but quickly become confusing themselves. Applications where predictions of user preference or intent are required have not so far been successful at all.*

*So anything beyond promulgation and simple interpretation seems problematic. Once more than a simple inference is attempted we seem to hit a brick wall. We realised this with the Badge system a decade ago. Interpreting the sighting of three or more Badges in a single space was presented as a “meeting”. However even in an office environment there are many reasons for three or more sightings at one place (meeting, tea time, passing in corridor). And that is before we extend to home or other environments.*

*One potential research direction is to provide much more feedback to the user. When we move a cursor on a screen it is clear where it is and what is likely to happen when we click. When walking through space it is much less obvious what the options are and how to control them. So visual and aural feedback with perhaps every nearby wall being used as a display may be one approach. The user might then be able to keep up as the context keeps changing. If proxy decisions are being made the reasoning can now be presented more easily. The user can interact in a much more informed way and help guide any decision-making process.*

*Perhaps a way to make progress beyond the engineering level is to imagine a “perfect” sensing system with full coverage of the environment. How would we define the context (world knowledge), semantics of queries, and user intent? How would the user interact to resolve ambiguities? Are statistical techniques likely to make useful predictions or are there too many plausible choices at each point? Could a series of*

*functional tests be devised which would give us the foundations to build on? It seems we are a long way from finding answers and only by moving away from unrealistic ambitions will we prevent the research area being discredited in due course.*

This is an abridged and updated version of the Royal Society Clifford Paterson Lecture 1999. The original paper was published in Phil. Trans. R. Soc. Lond., A, (2000), Volume 358, Pages 2349-2358, August 2000.