# A New Location Technique for the Active Office

Andy Ward[*], Alan Jones[†], Andy Hopper[*†]

Configuration of the computing and communications systems found at home and in the workplace is a complex task that currently requires the attention of the user. Recently, researchers have begun to examine computers that would autonomously change their functionality based on observations of who or what was around them. By determining their context, using input from sensor systems distributed throughout the environment, computing devices could personalize themselves to their current user, adapt their behaviour according to their location, or react to their surroundings. We present a novel sensor system, suitable for large-scale deployment in indoor environments, which allows the locations of people and equipment to be accurately determined. We also describe some of the context-aware applications that might make use of this fine-grain location information.

## Introduction

The modern home and office are equipped with sophisticated computing and communications devices, many of which require significant effort or specialist knowledge to configure and use effectively. Whilst the complexity of such devices will surely increase in the future, it may be possible to make them more user-friendly by transferring some of the configuration burden to the devices themselves. These computers would be *context-aware*, changing their behaviour based on how and where they were being used.

A context-aware computer or application must be able to determine the state of its surroundings. One method of discovering context is to monitor the locations of objects in the environment. In this paper, we first present an overview of research into location-aware computing and evaluate currently available location sensor technologies. We then describe a new location sensor, tailored to provide information for context-sensitive computers, which has been developed at the Olivetti and Oracle Research Laboratory (ORL). Finally, we examine potential applications of this system in an *Active Office* [1] where location-aware equipment will be commonplace.

## Location-aware Computing

Much of the existing research into context-aware computing has used location information provided by *Active Badges*[‡] [2][3], small computing devices worn by personnel. Each badge has a globally unique code that is periodically broadcast through an infrared interface. The infrared signals reflect off walls and furniture to flood the surrounding area, and are picked up by a network of sensors placed around the building. By determining which badges were seen by which sensors it is possible to deduce the location of a badge, giving a hint to the location of the badge's owner. Applications in which Active Badge information has been used include telephone call routing, security and environmental control [4].

---

[*] University of Cambridge Computer Laboratory, Pembroke Street, Cambridge, CB2 3QG, UK.

[†] Olivetti and Oracle Research Laboratory, 24a Trumpington Street, Cambridge, CB2 1QA, UK.

[‡] Active Badge is a registered trademark of Ing. C. Olivetti & C., S.p.A.

An extension to this system allows equipment to be tracked using a low-power version of the Badge called an *Equipment Tag* [1]. The developers describe a ‘nearest printer’ service offered to users of portable computers. Tags placed on the computer and printers report their positions, and the computer is automatically configured to use the nearest available printer as it is moved around a building.

The *ParcTab* [5] is a Personal Digital Assistant (PDA) that uses an infrared-based cellular network for communication. The infrared transmissions from ParcTabs can be used to determine their locations in the same way as Active Badges are located. Schilit *et al.* describe the use of the ParcTab system to implement applications involving context-triggered actions and automatic reconfiguration [6]. The ParcTab has also been used to implement a memory prosthesis [7] in which information about the user's context is collected and organized to form a biography that can be consulted at a later time.

Weiser has considered how the widespread deployment of location-aware devices might change the way we interact with computers [8]. He considers a vision of *Ubiquitous Computing*, in which computing elements are integrated into the environment to such an extent that they become invisible to common awareness. There will be a number of different types of device in this computing fabric, ranging in size to support different tasks. However, devices will not be specialized to a particular task—instead, they will be capable of adapting their behaviour based upon what is happening around them.

## Sensor Technologies

Systems like the Active Badge and the ParcTab are robust, relatively cheap, and can be integrated into everyday working environments. However, they locate objects only to the granularity of rooms, which act as natural containers for the infrared signals emitted by the devices. This limits the extent to which applications can adapt based on information from the system. It is therefore pertinent to consider other sensor technologies that might give finer-grain location information about objects in the office and home.

Electromagnetic trackers [9][10] can determine object locations and orientations to a high accuracy and resolution (around 1mm in position and 0.2° in orientation), but are expensive and require tethers to control units. Furthermore, electromagnetic trackers have a short range (generally only a few metres) and are sensitive to the presence of metallic objects.

Optical trackers are very robust, and can achieve levels of accuracy and resolution similar to those of electromagnetic tracking systems. However, they are most useful in well-constrained environments, and tend to be expensive and mechanically complex. Examples of this class of positioning device are a head tracker for augmented reality systems [11], and a laser-scanning system for tracking human body motion [12].

Radio positioning systems such as GPS and LORAN [13] are very successful in the wide-area, but are ineffective in buildings because of the reflections of radio signals that occur frequently in indoor environments. In-building radio positioning systems do exist (for example, the work of Feuerstein and Pratt [14]), but offer only modest location accuracies of around 50cm or more.

Location information can also be derived from analysis of data such as video images, as in the MIT *Smart Rooms* project [15]. Accurate object locations can be determined in this way using relatively cheap hardware, but large amounts of computer processing are required. Furthermore, current image analysis techniques can only deal with simple scenes in which extensive features are tracked, making them unsuitable for locating many objects in cluttered indoor environments.

After studying the currently available sensor technologies we concluded that none was well suited to the task of generating fine-grain location information for use in context-aware computing. Such a sensor would be accurate, reporting the positions of objects in three dimensions to within 15cm of their true locations. It would be scalable, both in the number of objects located and the area covered by the system, and would have a minimum of impact on the environment it was monitoring. We have undertaken work to develop a location system that meets these requirements.

## A New Location Technique

The ORL ultrasonic location system extends the work of Figueroa and Mahajan [16] and Doussis [17], who describe a system for mobile robot positioning. Measurements are made of times-of-flight of sound pulses from an ultrasonic transmitter to receivers placed at known positions around it. Transmitter-receiver distances can be calculated from the pulse transit times, from which, in turn, the transmitter's location is found by multilateration.

### Hardware

A small, wireless transmitter is attached to every object that is to be located. The devices, shown in Figure 1 consist of a microprocessor, a 418MHz radio transceiver, a Xilinx FPGA and a hemispherical array of five ultrasonic transducers. Each prototype mobile device has a unique 16-bit address, is powered by two lithium cells, and measures 100mm×60mm×20mm.

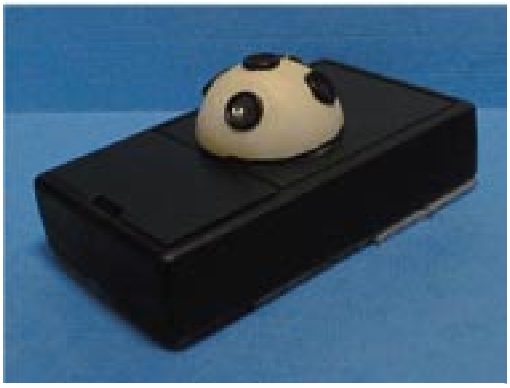

**Figure 1 - A mobile transmitter**

A matrix of receiver elements is mounted on the ceiling of the room to be instrumented. Receivers are placed in an array, 1.2m apart—the prototype system has 16 receivers in a four-by-four square grid. Each receiver has an ultrasonic detector, whose output is passed through an amplifier, rectifier and smoothing filter before being digitized at 20KHz by an ADC. The ADC is controlled by a Xilinx FPGA, which can monitor the digitized signal levels. Receivers also have a serial network interface, through which they are individually addressable, and are connected in a daisy-chain to a controlling PC.

Every 200ms, a radio message consisting of a preamble and 16-bit address is transmitted in the 418MHz band by a controller also connected to the PC. The PC dictates which address is sent in each message. The radio signals are picked up by the transceiver on each mobile device and decoded by the on-board FPGA. The single addressed device then drives its transducers for 50µs at 40KHz, and an ultrasonic pulse is broadcast in a roughly hemispherical pattern around the top of the unit. After receiving a message, mobile devices enter a power-saving state, activating themselves 195ms later, ready for the next message.

The controlling PC sends a reset signal to the receivers over the serial network at the same time as each radio message is broadcast. The FPGAs on each receiver then monitor the digitized signals from the ultrasonic detector for 20ms, calculating the moment at which the received signals peak for the first time[§]. The short width of the ultrasonic pulse ensures that receivers detect a sharp signal peak. The controlling PC then polls the receivers on the network, retrieving from them the time interval between the reset signal and detection of the first signal peak (if any signal was detected).

---

[§] Although the transmitter sends only one ultrasonic pulse, reflections of this pulse from objects in the environment may also reach the receivers, causing them to detect multiple signal peaks.

## Distance calculation

For each receiver, the interval $t_p$ between the start of the sampling window and the peak signal time represents the sum of several individual periods:

- $t_r$, the radio signal transit time from the controller to the addressed mobile device.
- $t_u$, the ultrasound transit time from transmitter to receiver.
- A number of fixed delays, $d_1 \ldots d_n$, such as the time taken for the FPGA to decode the radio message.

We then have

$$t_p = t_r + t_u + \sum_{i=1}^{n} d_i \quad (1)$$

Also

$$t_r = \frac{l_r}{c}$$

$$t_u = \frac{l_u}{v_s}$$

where $l_r$ is the distance from the controller to the addressed device, $c$ is the speed of light, $l_u$ is the transmitter-receiver distance and $v_s$ is the speed of sound in the room.

Since the controller and receiver matrix will normally be collocated, $l_r \sim l_u$. We also have $c \gg v_s$, so $t_u \gg t_r$, and we can therefore rearrange Equation 1 as

$$l_u \approx v_s \times \left( t_p - \sum_{i=1}^{n} d_i \right) \quad (2)$$

By empirically determining the fixed delays $d_1 \ldots d_n$ and making an estimate of $v_s$ based on the ambient temperature, we can use Equation 2 to calculate the transmitter-receiver distance from the time at which the first signal peak was detected.

## Position calculation

This principle of multilateration is demonstrated in Figure 2; a transmitter known to be a distance $x$ from a receiver must be located on a sphere of radius $x$ centred on that receiver. Four such spheres around receivers placed in three-dimensional space, such that they are not coplanar and no three are collinear, will intersect at only one point. The transmitter must have been located at this point in order to generate the observed distances.

In the ORL system all receivers lie in the plane of the ceiling, and the transmitters must be below the ceiling. This allows calculation of transmitter positions using only three distances, rather than the four required in the general case. Furthermore, we can use knowledge of additional distance measurements to refine our position estimates, making them less susceptible to errors in those measurements.

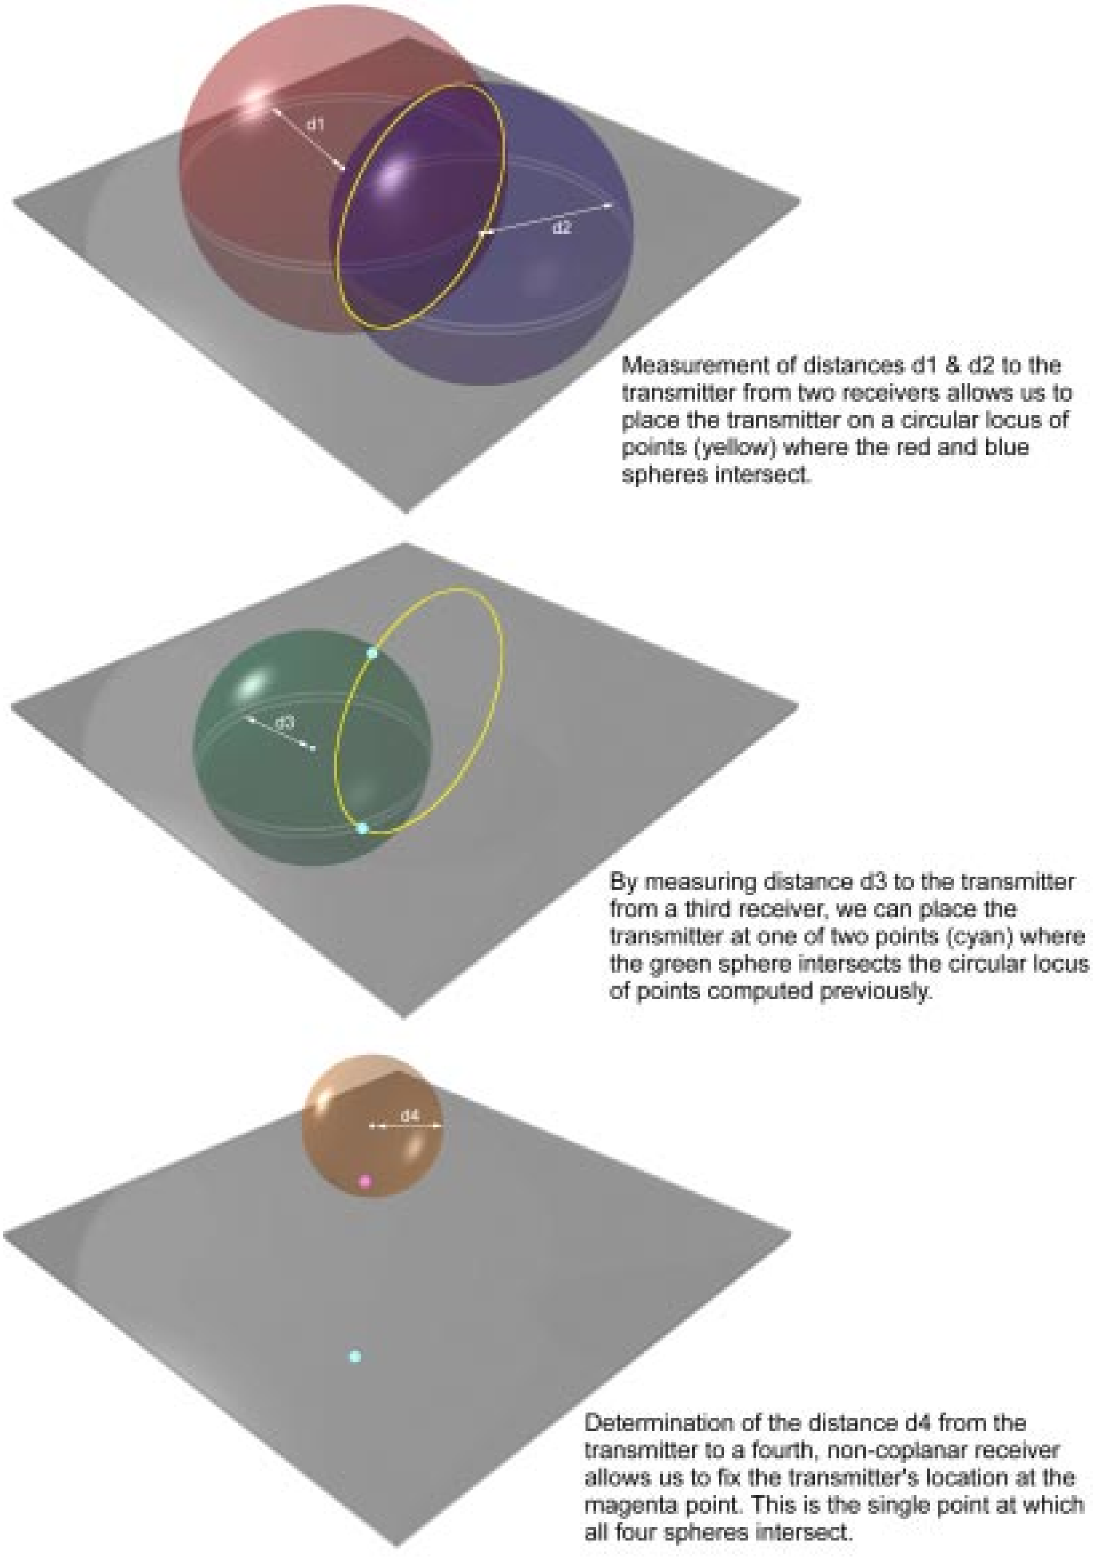

Measurement of distances d1 & d2 to the transmitter from two receivers allows us to place the transmitter on a circular locus of points (yellow) where the red and blue spheres intersect.

By measuring distance d3 to the transmitter from a third receiver, we can place the transmitter at one of two points (cyan) where the green sphere intersects the circular locus of points computed previously.

Determination of the distance d4 from the transmitter to a fourth, non-coplanar receiver allows us to fix the transmitter's location at the magenta point. This is the single point at which all four spheres intersect.

**Figure 2 - Position finding by multilateration**

Suppose a transmitter is at the coordinate (*u*,*v*,*w*), and the distance from it to a receiver at the coordinate (*x*,*y*,*0*)[**] is *l* . It can be shown that

$$l^2 = \left(x^2 + y^2\right) + \left(u^2 - 2xu\right) + \left(v^2 - 2yv\right) + w^2 \quad \textbf{(3)}$$

We can regard Equation 3 as a nonlinear model [18], and use nonlinear regression to fit the values of *l*,*x* and *y* for several receivers to this model. This gives estimates $\underline{u}$, $\underline{v}$ and $\underline{w}^2$ of the parameters *u*, *v* and $w^2$. We can then determine a best least-squares estimate for the transmitter position as the coordinate ($\underline{u}$,$\underline{v}$,-$\sqrt{\underline{w}^2}$), taking the negative square root of $\underline{w}^2$ to fix the transmitter below the ceiling[††]. The nonlinear model has three degrees of freedom, and knowledge of at least three transmitter-receiver distances is therefore required to calculate the transmitter position. Furthermore, the model cannot be fitted to the data if all receivers that detected a signal are collinear.

Reflected signals from objects in the environment can lead to incorrect distance measurements. Normally, the first signal peak detected by a receiver will be due to a pulse travelling along a direct line from the transmitter. This pulse will arrive before any reflected pulses, which must travel along longer paths. The distance thus measured by the system will be that of a straight line joining transmitter and receiver. Occasionally, however, the direct path may be blocked, and the first received signal peak will be due to a reflected pulse. In this case, the measured transmitter-receiver distance will be greater than the true distance, leading to an inaccurate estimate of the transmitter's position.

We have developed two techniques for identifying and eliminating inaccurate distance measurements. First we note that the difference of two transmitter-receiver distances cannot be greater than the distance between the receivers. If, by comparing pairs of measurements, we find two receivers whose results do not satisfy this test, we can state that the larger of the two distances must be a measurement along a reflected path (remembering that reflections can only increase the measured distance). We can then discard that result from the data set.

A second, statistical test is based upon the observation that the proportion of receivers that detect only reflected signals is small. Studentized residuals [19] provide one method of identifying outliers in data sets, and can be calculated for each of the distance measurements during the nonlinear regression process. An incorrect measurement will be considered to be an outlier in the full set of measurements, and it is likely to have a large studentized residual. We therefore remove the result with the largest positive studentized residual from the set of distance measurements (remembering, again, that reflections can only increase the measured distance), before recomputing the nonlinear regression and residuals.

The statistical test is repeated until the variance of the remaining measurements falls below an acceptable threshold (suggesting that all outlying data points have been eliminated), or only three measurements remain. A final calculation of the transmitter's position is then made using those data values. Simulations suggest that the two tests, when used together, can identify and eliminate all incorrect distance measurements caused by reflections in more than 90% of data sets. Furthermore, less than 4% of correct distance measurements are erroneously eliminated by the algorithms.

## Orientation calculation

Information from the ORL system can be used to find the orientation of objects. If three mobile devices are placed on a rigid body at known, non-collinear points, then by calculating the positions of the devices the orientation of the object can be found, as shown in Figure 3. This technique is most suitable for stationary or slow-moving objects, because the three devices will be located at different times—any intervening movement of the object will introduce errors into the calculated orientation. For the same reason, it is advantageous to locate the three devices consecutively. Orientations of objects whose motion is constrained in a known way may be found using fewer positions—a single device placed on a door, for example, could be used to determine whether it was open or closed.

---

[**] All receivers lie in the plane of the ceiling.

[††] A shadow solution ($\underline{u}$,$\underline{v}$,+$\sqrt{\underline{w}^2}$) corresponds to an impossible transmitter position above the ceiling.

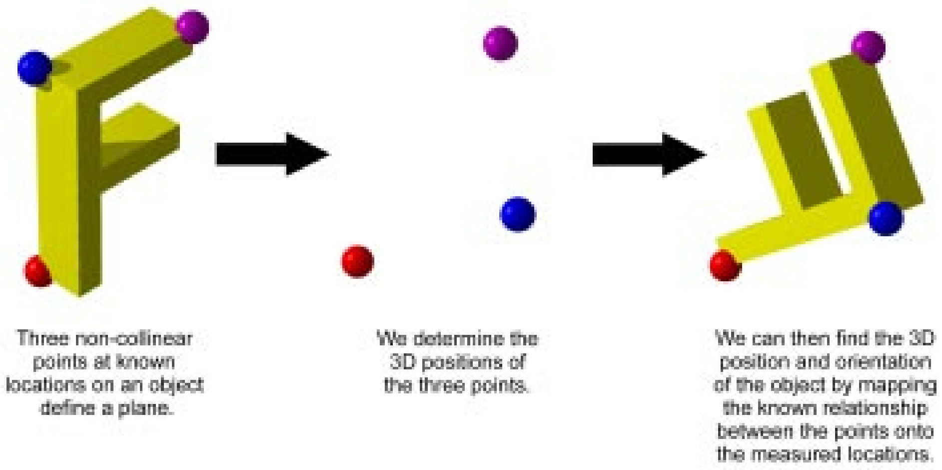

**Figure 3 - Orientation from location information**

Alternatively, a hint to an object's orientation can be obtained using knowledge of the set of receivers that detected ultrasonic signals from a mobile device. The directional transmission pattern of the ultrasonic pulse is known to be hemispherical, and some orientation information about a mobile device can therefore be found from its location and data on where the ultrasonic pulse was detected, as in Figure 4. If the device is rigidly affixed to an object at a known point and in a known orientation, an approximation for the object's orientation can then be deduced. This approach is most suitable when tagging of the object with multiple transmitters would be cumbersome.

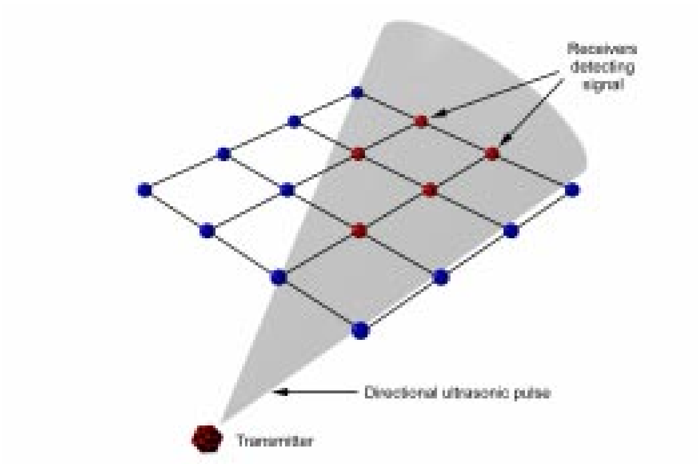

**Figure 4 - Single transmitter orientation**

Assessment

The prototype ORL location system, equipped with 16 ceiling receivers, operates over a volume of some 75m$^3$. Figure 5 shows the accuracy of the location information provided by the system. 95% of raw readings lie within 14cm of the true position, and a similar proportion of averaged readings (calculated as the mean position over ten cycles) lie within 8cm of the true position. The main factors limiting system accuracy appear to be the finite size of the transmitter array and noise in the receivers. The location information can be displayed in real-time within a VRML model of the ORL building, as shown in Figure 6, allowing simple interpretation of the data.

Up to $2^{16}$ mobile devices can be located by the prototype system, a number which could easily be increased by changing the size of the address space to, say, 48 bits. The addressable nature of the mobile devices allows flexible operation of the location system—different location qualities-of-service

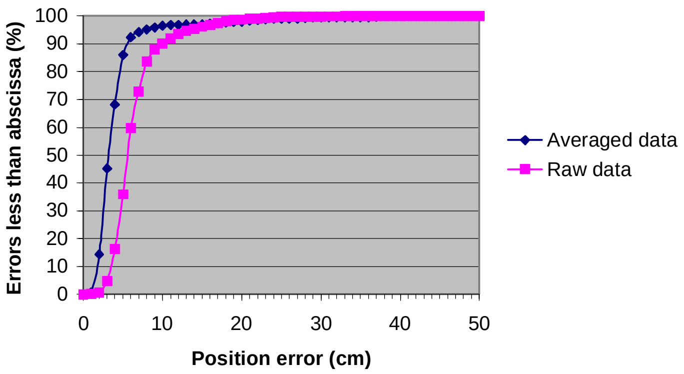

**Figure 5 - Position error cumulative probability distribution for ultrasonic location system**

can be allocated to different objects, reflecting the levels of demand for information about them. For example, a system that forwarded incoming telephone calls to the handset nearest the callee would require more location information about people, who are highly mobile, than about telephones, which are relatively static. The current system provides five location updates each second, which may be shared between objects in this way. Future implementations may support significantly higher update rates—tests have shown that a fast Pentium[‡‡]-class PC can perform at least 25 position computations each second, and the fundamental limit to the update rate is approximately 50Hz[§§].

The daisy-chain serial network can support up to 256 receivers. The largest volume over which the current system could operate is around 1100m$^3$, a figure that compares very favourably with the working volumes of other location sensors. At least three transmitter-receiver distances must be measured to compute the transmitter's location, but this requirement was found not to be restrictive, because of the large number of ceiling receivers and broad ultrasound transmission pattern from mobile devices.

**Figure 6 - VRML view of location and orientation data**

---

[‡‡] Pentium is a trademark of Intel Corporation.

[§§] Reverberations of the ultrasonic pulse in a typical office take up to 20ms to die out, and this dictates the minimum interval between transmissions from mobile devices.

The system appears to work well in the office environment. Day-to-day maintenance is minimized by the power-saving technique used in the mobile units, which limits their average current consumption to 400µA and gives them a battery lifetime of around three months. No interference was seen from sources of background ultrasound in the office (such as cooling fans), which seem to be too quiet to affect the operation of the receivers. The highest sound pressure level recorded around the mobile devices was 110dB, within suggested exposure limits to ultrasound in the workplace [20].

## Applications

We intend to use information from the ORL system to improve a number of existing applications. Those that use Active Badge data to select the nearest telephone or printer to a particular person could make more accurate choices based upon the detailed location information. The *Teleporting System* [21], which can redirect an X Window System[***] environment to different computer displays, could use location and orientation data to present a user's familiar desktop on a screen that faced them whenever they entered a new location. Enhancements might also be made to videophones, which could automatically configure themselves by selecting suitable cameras, displays, speakers and microphones based on where users were and what input and output devices they could interact with.

Other types of application require fine-grain location information that cannot be provided by technologies such as the Active Badge. It would be possible to manage data streams between input and output devices based on their physical locations—a video stream could be set up between a camera and a display simply by touching the two together. Computer systems could be built from component parts placed near each other; an example might be a large display made from an array of smaller screens. Location devices might be used to signal gestures for controlling computers, or as an input tool that could pick information up from a display and place it on another screen or printer.

Fitzmaurice has described three-dimensional information areas around physical objects called *Spatially Situated Information Spaces* [22]. Users interact with these computer-synthesized spaces through location-aware PDA ‘portholes’; for example, a PDA placed near a printer might display the contents of the job queue. In the original research, the PDAs were located using tethered electromagnetic trackers—obviously, the use of a wireless location device such as that described in this paper would ease interaction with such spaces.

A wireless tracker could also replace a variety of other sensors. Hodges and Louie describe an *Interactive Office* [23] that gathers information about the activity of the occupants. This data is generated by a number of different sensor types, such as motion detectors and reed switches (which monitor movement of people and the positions of doors). The ORL location system could provide much of this information using a single, low-powered and untethered device, thus simplifying the physical and computing infrastructure required to support the interactive environment.

## Conclusions

We have demonstrated a system that can determine the location and orientation of objects within a building. The information provided by the system is sufficiently fine-grained to allow investigation of a new set of context-aware applications, and the system has a very large working volume. Furthermore, the wireless, low-powered nature of the location sensors allows them to be integrated into an everyday working environment with relative ease.

In the near future, we intend to extend the area covered by the prototype system and will investigate minaturization of the mobile devices. We will also conduct research into software architectures that can support a large-scale system deployment in which hundreds, perhaps thousands, of objects are located within a building.

---

[***] The X Window System is a trademark of The X Consortium.

# Acknowledgements

Andy Ward would like to thank the UK EPSRC and the Olivetti and Oracle Research Laboratory, Cambridge for their support of this work by means of a CASE award.

# References

[1] Harter, A., Hopper, A. *A Distributed Location System for the Active Office*.
IEEE Network, Special Issue on Distributed Systems for Telecommunications, January 1994.

[2] Want, R., Hopper, A., Falcão, V., Gibbons, J. *The Active Badge Location System*.
ACM Transactions on Information Systems, January 1992.

[3] Want, R., Hopper, A. *Active Badges and Personal Interactive Computing Objects*.
IEEE Transactions on Consumer Electronics, February 1992.

[4] Elrod, S., Hall, G., Costanza, R., Dixon, M., Rivieres, J. *Responsive Office Environments*.
Communications of the ACM, July 1993.

[5] Adams, N., Gold, R., Schilit, B., Tso, M., Want, R. *An Infrared Network for Mobile Computers*.
Proceedings of the USENIX Mobile & Location-Independent Computing Symposium,
Cambridge Massachusetts, August 2–3 1993. pp. 41–51.

[6] Schilit, B., Adams, N., Want, R. *Context-Aware Computing Applications*.
Workshop on Mobile Computing Systems and Applications, Santa Cruz, December 1994.

[7] Lamming, M., Flynn, M. *"Forget-me-not"—Intimate Computing in Support of Human Memory*.
Proceedings of FRIEND21, International Symposium on Next Generation Human Interface,
Meguro Gajoen, Japan, 1994.

[8] Weiser, M. *The Computer for the 21st Century*.
Scientific American, September 1991.

[9] Raab, F., Blood, E., Steiner, T., Jones, H. *Magnetic Position and Orientation Tracking System*.
IEEE Transactions on Aerospace and Electronic Systems, Vol. AES-15, No. 5, September 1979.

[10] Ascension Technology Corp. *'Flock of Birds' Technical Description*.
Burlington, Vermont, 1994.

[11] Wang, J., Chi, V., Fuchs, H. *A Real-time Optical 3D Tracker for Head-mounted Display Systems*.
Computer Graphics, Publication of the ACM SIGGRAPH, Vol. 24, No. 2, 1990. pp. 205–215.

[12] Sorensen, B., Donath, M., Yang, G., Starr, R. *The Minnesota Scanner: A Prototype Sensor for Three-Dimensional Tracking of Moving Body Segments*.
IEEE Transactions on Robotics and Automation, Vol. 5, No. 4, August 1989.

[13] Sonnenberg, G. *Radar and Electronic Navigation*.
Butterworths, 1988.

[14] Feuerstein, M., Pratt, T. *A Local Area Position Location System*.
IEE Conference Publication No. 315, 1989. pp. 79–83.

[15] Pentland, A. *Machine Understanding of Human Action*.
Proceedings of 7th International Forum on Frontier of Telecommunication Technology, Tokyo, Japan, November 1995.

[16] Figueroa, F., Mahajan, A. *A Robust Navigation System for Autonomous Vehicles using Ultrasonics*.
Control Engineering Practice, Vol. 2, No. 1, 1994. pp. 49–59.

---

[17] Doussis, E. *An Ultrasonic Position Detecting System for Motion Tracking in Three Dimensions*.
PhD Thesis. Tulane University. July 1993.

[18] Myers, R. *Classical and Modern Regression with Applications*.
PWS-KENT, 1990.

[19] Glantz, S., Slinker, B. *Primer of Applied Regression and Analysis of Variance*.
McGraw-Hill,1990.

[20] Cracknell, A. *Ultrasonics*.
Wykeham, 1980. pp. 186–190.

[21] Richardson, T., Mapp, G., Bennett, F., Hopper, A. *Teleporting in an X Window System Environment*.
IEEE Personal Communications Magazine, Vol. 1, No. 3, Third Quarter 1994.

[22] Fitzmaurice, G. *Situated Information Spaces and Spatially Aware Palmtop Computers*.
Communications of the ACM, July 1993.

[23] Hodges, S., Louie, G. *Towards the Interactive Office*.
Proceedings of SIGCHI'94, Boston, April 1994.