that implements the projector's user interface. From there she can advance slides as required, controlling the shared projector with her personal PDA.

The versatility of the web model of interaction, including CGI and client upload via forms, can be exploited for less obvious applications. A Cooltown camera recognizes the user by acquiring her URL (which contains a special subsection dedicated to web clients that are cameras), perhaps through a user badge (see section 2.5.1). Then, by surfing to that section, the camera finds a form through which photos can be uploaded to the user's page via HTTP POST.

Huang *et al.* [135] interestingly describe the practical challenges faced by a non-technical user wishing to publish digital photographs on his web page: even when there is wireless connectivity between the camera and the computer, the operation is far from trivial and typically involves tedious repetitive operations such as the renaming of dozens of files. Ubicomp environments will live up to their promise only after solving all these practical issues. Cooltown's choice of web protocols is a stable and well-understood base for experimentation in this domain.

2.5 ORL/AT&T Labs Cambridge

The Olivetti Research laboratory in Cambridge, UK (later Olivetti-Oracle Research Laboratory, now AT&T Laboratories Cambridge), where the Active Badge was invented in 1989, was an early player in the ubicomp game. This systems-oriented research centre developed several new technologies revolving around the global vision of **sentient computing** put forward by the lab's director Andy Hopper.

The core idea of sentient computing is to make computing systems reactive to the physical world: as the user's context changes, the applications should adapt to the new environment and serve the user in the manner that is most appropriate to the new circumstances. The enabling technology that makes this possible is that of inexpensive digital sensors (including cameras, microphones and compasses) that could be deployed in large numbers. Sentient computing leads to "do-nothing technology" (the user's presence is sufficient to make things happen, as opposed to having to request things explicitly with button presses) and to "global smart personalization" (systems automatically reconfigure themselves with the preferences of the current user). Says Hopper [134]:

> Central to any good working relationship is a degree of mutual understanding between the two parties. The problem with conventional human-computer interaction is that responsibility for understanding, or the lack of it, lies wholly with the user. [...] Instead of bringing the user to the computer, let us take into account that people live and work in a real physical world, and make this notion—the concept of space—integral to the operation of our computer systems. We need to make computer systems aware of the physical

> environment—shifting part of the onus of understanding from user to machine. Awareness comes through sensing, and that implies the need for appropriate sensor technologies to collect status and location data.

Many research groups around the world are currently working on the related topic of “context-aware computing”—systems that reconfigure themselves depending on the “context” (whatever this may mean) of their user. It is interesting to note the significant number of cases in which the context information consists primarily of the user’s *location*.

Outdoors, the Global Positioning System provides absolute geographical coordinates with an accuracy of around 10 m. Originally developed for military uses, the technology was found to be extremely useful for civilian air and sea navigation and is now being widely deployed in commercial car navigation systems[10]. Handheld receivers are also available for bikers and trekkers.

The ORL/AT&T laboratory has been investigating location technologies for over a decade and has been responsible for pioneering contributions in the field of indoors location systems—from the Active Badge, whose spatial resolution is at the room scale, to the Active Bat, accurate to a few centimetres; and from the vision-based TRIP system, which locates objects via special 2D bar codes, to the weight-sensitive Active Floor, which can recognize people from their walking patterns.

The ubicomp-related research conducted at AT&T Labs Cambridge has not been limited to location technologies, though: further relevant projects have included Virtual Network Computing (VNC), which lets you access your personal computing environment from anywhere, and the Prototype Embedded Network (PEN), a short-range wireless communication system with an emphasis on low power.

2.5.1 The Active Badge

The Active Badge [254], in a pattern common to many other inventions, was born out of frustration. Our workplace was spread out on three floors of a narrow building and the researchers would frequently be in places other than their own offices, discussing ideas with colleagues; finding a specific person often involved phoning their office and then, in sequence, the hardware lab, the meeting room, and the offices of other colleagues who might know the whereabouts of the intended callee. One of the researchers, Roy Want, took up Hopper’s half-serious challenge to create a system that addressed this problem, and came up with the Active Badge (figure 2.7 shows a later version of the device, *circa* 1992).

[10] At the time of writing (2001) these are mostly seen in high-end rental cars in Europe and America, but they are commonplace consumer electronic items for motorists in Japan.

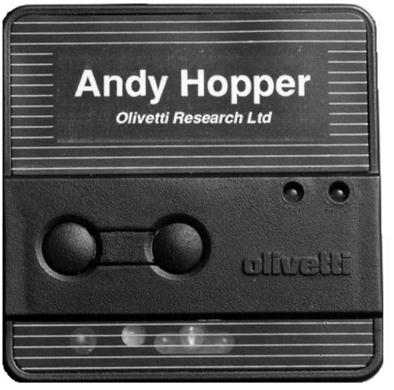

Figure 2.7. The Active Badge (courtesy of AT&T Laboratories Cambridge).

The principle on which the Active Badge is based is very simple. The device is worn by users like a conventional name badge—which it also is. Its infrared transmitter periodically sends out the badge's identifier, which is picked up by fixed sensors deployed throughout the building. Because the infrared signal bounces off walls without penetrating them, the detection of a signal by any given sensor indicates that the corresponding badge is located in the same room as the sensor. By collating sightings from all the sensors in the building, the system knows where individual badges are and can offer services such as the automatic redirection of phone calls. The hardware is cheap and easy to build because the underlying technology of infrared signalling is widespread in consumer equipment such as remote control units and is therefore commercially mature, with its components widely available as commodity items.

This system was first deployed across the lab as an experimental prototype in early 1990 and was subsequently adopted by other academic and industrial research institutions in Europe and America including the University of Cambridge, Imperial College London, the University of Twente, Digital Equipment Corporation, Xerox, Bellcore and MIT. The Badge was at one point even sold as a product by our parent company Olivetti, though it never reached commercial acceptance in the marketplace. We used the system continuously for about ten years until we replaced it with the higher resolution Active Bat system (described in section 2.5.3 next); this means that we can now speak about it from extensive practical experience and not just as an intellectual experiment.

In our lab, the Active Badge soon became an essential part of the infrastructure: researchers really liked the fact that they could freely wander away from their office even if they were expecting an important phone call[11]. Sometimes, when a new start-up company was formed from one of our research projects, their move to new premises was followed by the deployment of an Active Badge network, because the researchers from our lab hated to forgo the convenience of the badge after having experienced it at ORL.

Wearing the badge was always voluntary, but few if any people ever decided to opt out. I did, for a certain time, in part also as a kind of social service to ensure that new employees would not be forced to wear one because of unanimous peer pressure; but this did not change the overall picture, and most staff members have been consistently viewing the system as a sci-fi asset of the lab rather than as a privacy threat. The following statement by Want *et al.* [254] (emphasis mine) has been validated by years of daily use:

> There is a danger that in the future this technology will be abused by unscrupulous employers. If it is, then legislation must be used to protect us from this kind of technology abuse. However, it is our conclusion that *amongst professional people responsible for their own work time*, it is a very useful and welcome office system.

This acceptance was fostered by some explicit policy decisions: firstly, badge sightings were never logged—they were used instantaneously and then thrown away. Secondly, we committed to **reciprocity**. Everybody wore a badge, and everybody had the `xab` application (figure 2.8) running on their workstation all the time; this displayed a list of all the staff members with their current location, nearest phone number and number of people in the same room. So, yes, your boss could watch you, but in return you could watch him yourself if you were so inclined. As a matter of fact the boss even installed badge sensors in his *home*, which relayed his position via ISDN to the badge server at the lab[12]. This “ethical foundation” of reciprocity helped make it clear that the purpose of the system was not to monitor staff but to help lab members find each other.

At some point a new facility was added to the system: you could now set a “watch” on someone and you would get a beeping alert when their location changed. This was meant for the situation in which *A* wanted to contact *B*, but the badge system told *A* that *B* was in a meeting, which *A* did not want to disturb. By putting a watch on *B*, *A* could catch him just as he walked out of his meeting. One of the researchers complained that this would facilitate abuse of the system by curious snoops or petty thieves who could sneak into someone’s office during the lunch hour, safe in the knowledge that they would get advance warning of their

[11] This was before everyone in the civilized world (except yours truly who still stubbornly refuses to carry one) had a cellular phone.

[12] He did this primarily to support automatic *teleporting* of his desktop session between work and home, described next on page 33.

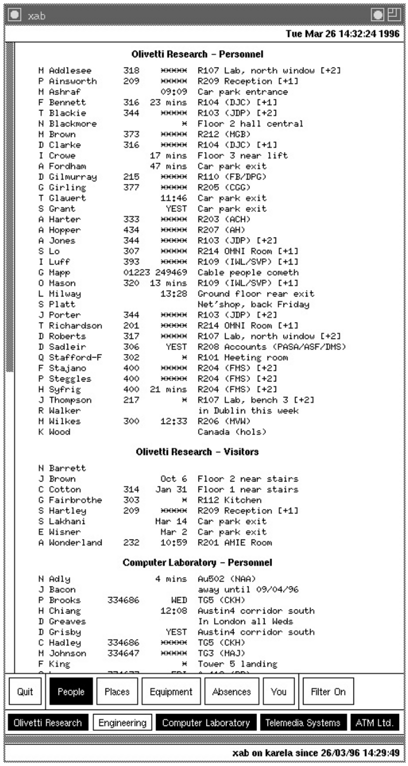

Figure 2.8. The `xab` application, displaying badge sightings on an X Windows workstation (courtesy of AT&T Laboratories Cambridge).

victim's return. The issue was addressed by resorting to the reciprocity principle: whenever user *A* set a watch on user *B*, *B* would receive a notification. This made the feature acceptable and a new social protocol naturally developed around it: on receiving the notification, *B* would phone *A* to say "I'm back now, were you looking for me?".

Many other projects in our lab made use of the location information provided by the Active Badge to enhance some other service. In our distributed multimedia systems Pandora [133] and Medusa [265] you could videophone someone by name as well as by phone number[13], and videomail would be automatically annotated

[13] Observe that "by phone number" effectively means "by location" unless we are dealing with mobile telephony: when you dial 321123 you are not actually calling "George", but "George's office". Cellular phones reverse this situation—and therefore are wide-area person-tracking devices in their own right—in part also owing to the interesting social convention that the mobile terminal is a strictly personal item that is never shared. When the fixed phone on your desk rings and you're not in the immediate vicinity, it is acceptable for a member of your family (at home) or a colleague (at work)

with the names of the people present in the room at the time of the recording. We also experimented with follow-me audio and video, where media streams would be redirected while you moved from one office to another.

Figure 2.9. The Smart Beverage Dispenser (courtesy of AT&T Laboratories Cambridge).

Ward, Naylon *et al.* interfaced the coffee machine[14] to the badge system so that clicking the badge button when walking up to the machine selected the user's preferred beverage from a list featuring, among others, coffee, chocolate "and the revolting *mokaccino*, a subtle blend of coffee and chocolate" [24]. This can be seen as a light-hearted example of "global smart personalization": the appliance (see figure 2.9) automatically adapts its behaviour to that preferred by its current user.

Another example of global smart personalization comes from the teleporting

to answer it for you; but it is considered bad manners for anyone to do this with your *mobile* phone, even when you temporarily leave it on your desk.

[14]It is interesting to note the frequency with which hackers take caffeine not just as a substance that keeps them awake during late-night programming sessions but also as an inspiration for whacky computer recreations. See for example the CMU coke machine that users could query via `finger(1)` to ensure that they would be issued a can that had had time to cool down [58]; the Trojan Room coffee pot at the University of Cambridge [235], later to become the first webcam; the ParcTab-augmented coffee pot of Xerox PARC, which allowed a public-spirited researcher to alert colleagues electronically whenever a fresh pot had been brewed [258]; and the Hyper Text Coffee Pot Control Protocol (HTCPCP/1.0), which made its way to an appropriately dated Internet RFC [184].

system [217] and from its better known open-source successor VNC [218, 25] (Virtual Network Computer), both developed by Richardson *et al.*

These systems allow a user to maintain a GUI session (a “desktop”) that can be accessed from any other computer used as a thin client, with the state of all open applications being preserved as the user moves from one thin client to the next. They implement personalization in the sense that any computer can become “your” computer just by connecting to the networked server that runs your desktop. You can use any currently “physically free” computer (meaning one with nobody sitting in front of it) to access your remote resources, even one whose screen is locked or taken by someone else’s session. You don’t have to log in—the thin client that will display your remote desktop is overlaid on top of whatever is already there, without disturbing the existing session.

This Virtual Network Computer paradigm can be very useful on its own but becomes even better if augmented by location information: you can then simply walk up to a computer anywhere and have “your” desktop appear on it, with the cursor still in the middle of the sentence you were typing, and without having manually to launch the VNC client and feed it your password. Note that in actual practice the “teleporting” action (i.e. making the remote desktop appear on the local workstation) is not triggered automatically by the movement of the user: we also require an explicit click of the badge button. It would otherwise be embarrassing to have one’s desktop randomly appearing on nearby screens—particularly ones being used by others—as one moved around the building. The click action also allows the user to select a specific destination screen in rooms where several are present. Having said that, the action of teleporting *away* (to nowhere) is indeed automatic, and triggered simply by the user moving to another location. This do-nothing operation is also an advantage for security: since your desktop disappears from the screen once you leave the room, there is no danger of accidentally leaving your session open for a passer-by to abuse.

In the above we may observe the badge being used as an authentication device. Because the badge is considered to be a personal device that is always with its owner, it is treated like a physical key. Whoever holds my key can open my filing cabinet, and whoever holds my badge can recall and access my VNC desktop. When the badge started being used in this way (which happened when we allowed it to unlock doors, way before teleporting) we introduced a challenge-response protocol to thwart replay attacks (see section 3.4.3).

It is interesting to note that by far the most commonly used function of the Active Badge system turned out to be not the automation of some manual task but simply the ability to view, in a very basic textual form (see figure 2.8 above), the location of all the lab members. Most people at the lab always keep a badge window running for the whole duration of their login session. Apart from the obvious use of figuring out where to go or phone to actually meet a specific colleague, a more subtle social benefit is the ability to check in advance whether the colleague is free or already busy with someone else.

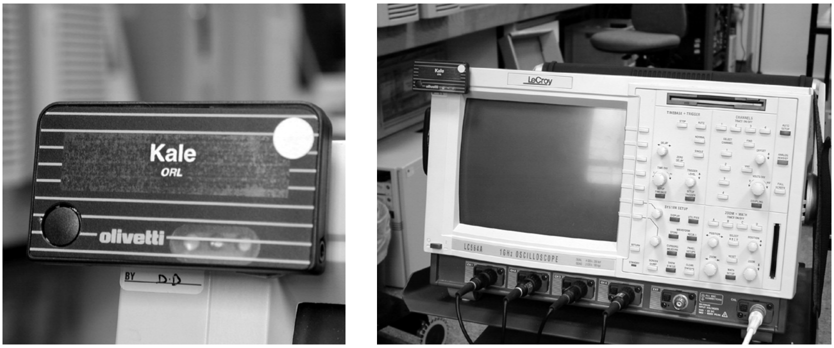

Figure 2.10. A tob (“transmit-only badge”) is used to track the location of this trolley-mounted oscilloscope (courtesy of AT&T Laboratories Cambridge).

We soon realized that the facility to locate colleagues could be fruitfully extended to that of locating frequently borrowed objects. Prime candidates for this were the “good” oscilloscope (this was mounted on a trolley so that it could be wheeled from the main hardware lab to wherever it was needed; but engineers would frequently hoard it in their offices for days, to the dismay of their colleagues who had to broadcast distressed-sounding email requests; see figure 2.10) and the label maker whose successor is shown in figure 5.1 on page 115. Originally this was done by sticking ordinary badges on the objects; in 1993 we developed a smaller, simpler and more energy-efficient “transmit-only[15] badge”, or tob, which we ended up deploying on all our computer equipment—printers, computers, switches, cameras, microphones, displays and other networked peripherals.

While the usefulness of tracking the location of our peripatetic oscilloscope is immediately obvious, it may seem odd to want to label items such as printers or desktop computers which, after all, don’t move very often. The benefit of this practice is that it enables an application running on a given machine to ask the system what equipment is available in the same room. Relying on static configuration files for the same purpose would be a maintenance nightmare, especially in the experimental environment of a research lab where researchers continually augment their

[15] While a transmission facility was obviously necessary for any badge to communicate its location, the full version of the badge could also *receive* messages from the network. This allowed additional functionality such as paging (the arrival of the sandwich man prompted the receptionist to issue a broadcast that made all badges play a tune recognized by users as “food, glorious food!”) and challenge-response.

workstations with new gadgets and *do* move equipment around more often than you might think. In fact the badge-driven teleporting described earlier in this section relies heavily on tob information to find out which machines in the current room can work as a VNC viewer.

2.5.2 The Active Floor

The Active Badge requires you to tag every object or person you wish to track. We also decided to investigate systems that did not impose this requirement. A pressure-sensitive floor, for example (a staple of sci-fi action movies), can detect users walking on it. In the transition from science fiction to engineering, though, one has to address a variety of practical concerns.

Firstly, what should be the spatial resolution of the sensing floor? If too low, it will be sufficient to activate a *Mission Impossible*-style alarm but not to detect where an individual is in the room. If too high, cost will become an issue, and it will be difficult even to route wires from the individual sensors to the electronic circuits that process and aggregate their readings.

Secondly, the requirements for mechanical robustness of the sensors are not trivial to fulfil. A floor must be able to withstand fairly rough treatment above and beyond being walked upon—including having heavy objects dropped upon, being subjected to shearing stress when someone runs and stops, being subjected to high pressure from canes, high heels or chair legs, having liquids spilled on and so on. Furthermore, the fact that the floor is two-dimensional obviously means that the number of (suitably rugged) sensing elements to be deployed grows with the square of the chosen spatial resolution. One proposal of adopting boolean microswitches arranged in a sub-centimetre grid (in the hope of getting a simple bitmapped output that would show the shape and orientation of the shoes of the people treading on it) attracted the comment that it would be cheaper to pave the floor with banknotes of the highest denomination.

The final design, described by Addlesee *et al.* in [5], consists of 0.5 m $\times$ 0.5 m carpeted floor tiles made of plywood and steel, supported at the corners by load cells in the shape of cylindrical metal pads (see figure 2.11). Each cell supports four tiles and acts like a precision scale[16]. Compared to alternatives such as the microswitch array, this arrangement offers several practical advantages: the surface is sufficiently robust and stable to be used as an ordinary floor, installation is relatively easy and the tiles can be carpeted without affecting the floor's sensing ability.

[16]The resolution is about 50 g with a range of 0–500 kg and a precision of 0.5%. In the prototype Active Floor, the output of each cell is sampled at 500 Hz.

Figure 2.11. A detail of the construction of the Active Floor, showing a bare tile, a carpeted tile and a load cell (courtesy of AT&T Laboratories Cambridge).

As a sensing system, the Active Floor offers some interesting properties. First and foremost, it does not require prior marking of the items to be tracked—it will always accurately report the total weight of the objects in a given area. A peculiar feature that few other sensing systems possess is that, with the Active Floor, an object can never "obscure" another: to evade detection, items would have to be suspended from the ceiling. If any sizable object is moved from one desk to another (or brought in or taken out of that office), or if a person enters or exits the room, the floor will notice.

At the same time, though, it is difficult to extract higher level meaning from the raw sensor data. One can track the position of the centre of gravity of a moving object but it is difficult to establish the identity of the particular object that is moving. A comparison with a database of objects of known weights somehow brings us back to a priori tagging—and in any case the method is incapable of distinguishing multiple instances of the same artefact, e.g. three laptops of the same model. Besides, when several objects move at the same time, the problem of locating their centre of gravity no longer has a unique solution. Various heuristics may be applied, but a really robust solution is needed before the floor can reliably be used as a location system.

It is nonetheless possible to extract some pretty remarkable results out of the data produced by the floor: even though the experimental setup was somewhat constrained[17], the cited authors reported reasonable success in identifying people based

[17] A single subject was examined at a time, and the subject would cooperatively step in the middle

on their gait [5] thanks to HMM-based analysis techniques originally developed for speech and face recognition.

2.5.3 The Active Bat

The badge system could not resolve position at higher granularity than the room; as such it could not support some more advanced location-aware applications. For example the cited location-based VNC teleporting (section 2.5.1) works well with individual offices where each room only contains a few workstations, but it would be useless in an open-plan environment with tens or hundreds of desks all labelled as being in the same "location".

Starting in 1995, Ward *et al.* [257, 256] developed a new active tag that signalled its position using ultrasound as opposed to infrared. The name "Active Bat" for this successor of the badge is of course a reference to the flying mammal that uses a location system based on similar sensing technology.

Given the much lower speed of the ultrasonic signal compared to that of the infrared (think speed of sound vs. speed of light), it is possible to measure the time of flight of the pulse from transmitter to receiver with sufficient accuracy to yield a precise distance measurement between the two.

In the Active Bat system (figure 2.12), the ultrasound transmitter ("bat") is attached to the mobile object as the tag, while the receivers are placed at known locations on the ceiling. Bats are individually addressable and, at regular intervals, a base station sends a radio command (whose delivery is instantaneous compared to that of the ultrasound message) to a specific bat, ordering it to transmit an ultrasonic pulse. At the same time the base station notifies the ceiling receivers that that particular bat is now going to transmit. The receivers, laid out on the ceiling at regular intervals, measure different flight times depending on their relative location to the transmitting tag. Then, using trilateration[18], the system works out the position of the tag in three dimensions with respect to the known positions of the ceiling receivers. The controller then addresses another bat and starts a new cycle.

Substantial engineering has brought the initial prototype to a mature system that scales well both in the size of the area to be covered and in the number of bats that can be tracked simultaneously. The current version is deployed on all three floors of our laboratory and can locate bats with an accuracy of about 3 cm. The bats themselves are about the size of a thick lighter and have now completely replaced badges and tobs in our lab.

The software back-end that supports the Active Bat is considerably more elaborate than the one for the badge. Application programs would have little use for

---

of the designated target tile.

[18]This is the dual technique to the better known "triangulation". It uses the sides of a triangle, rather than its angles, to determine the position of the third vertex.

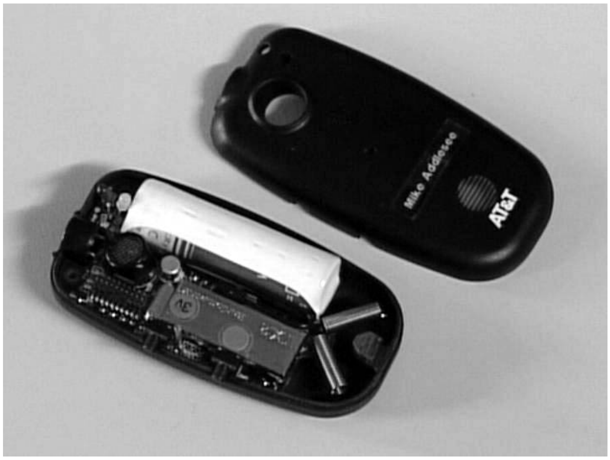

Figure 2.12. The Active Bat (courtesy of AT&T Laboratories Cambridge).

the actual centimetre-resolution 3D coordinates of the individual bats. A layer of spatial indexing middleware, developed by Steggles *et al.* [243], raises the level of abstraction by processing the bat sightings and generating events when something of interest occurs, such as when the containment zone of a tagged entity (say, a person) overlaps that of another (say, a wall-mounted screen, meaning that that person is now in a position to interact with the screen, point at parts of it and so on). Applications can register interest in specific spatial events and will be notified by the system when these occur.

High resolution location information enables users to interact in new ways with the ubiquitous computing system: you can activate individual objects in your surroundings, or signal your intention to use them, merely by going near them. But to what does this translate in practice? Perhaps one of the most interesting new uses for this enhanced information is the ability to define "virtual buttons". As a developer of ubiquitous computing applications you may print a label saying "email me an image of this whiteboard", paste it on the edge of the whiteboard, define a 3D region around the label, register it in the system and associate it with the appropriate sequence of actions. Your users can now click their Active Bat on top of that label and receive a dump of the whiteboard in their email inbox. (And note how the interpretation of "email *me*" follows the identity of the owner of the bat.)

As a programmer, you can define virtual buttons wherever you like in 3D space, which users can click using their bat as if it were a mouse. In our laboratory we now have distinctive "smart posters" on the walls (figure 2.13) that not only describe an application but also let you control it by clicking your bat over the appropriate self-describing areas of the poster [26, 23].

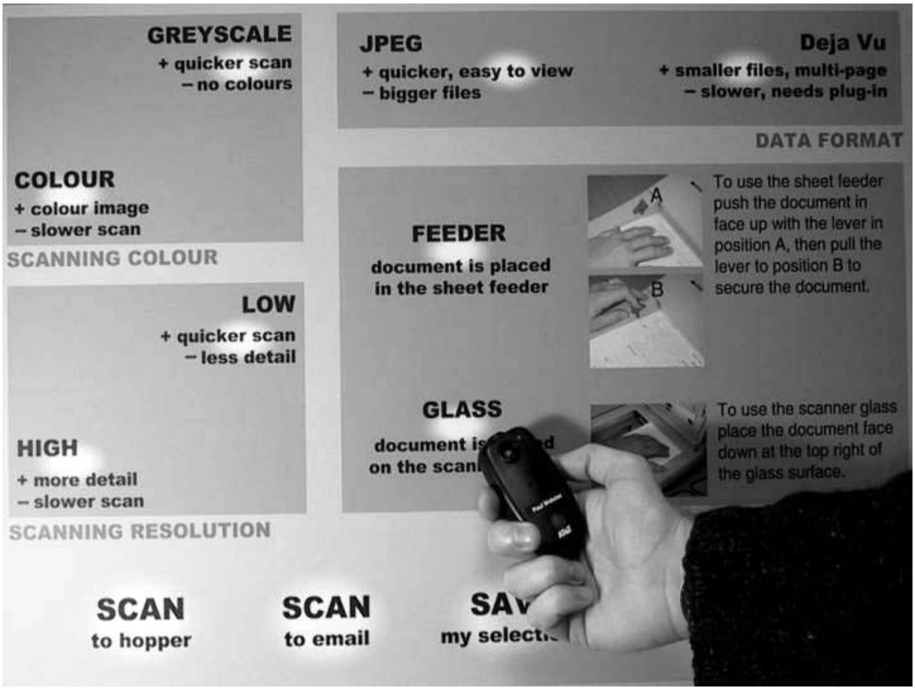

Figure 2.13. A smart poster acting as the computerless user interface for a public document scanner nearby. The user is about to scan the document currently on the glass (courtesy of AT&T Laboratories Cambridge).

Of course, having such buttons near walls makes it easier to label them with printed notices, but this is just a notational convenience. You can also define "activation regions" whose associated action is triggered simply by spatial overlap, without requiring a button click. For example you can equip your office with many cameras and have a follow-me videophone by associating an active region with the coverage area of each camera. As you walk around, from your desk to the whiteboard and from the window to the office of your colleague, the system automatically keeps you in shot by switching to the camera and microphone that best cover your current position.

Badge emulation (room-level location), desk-resolution teleporting, smart posters and follow-me video do not exhaust the uses of the Active Bat. Another interesting ubicomp development is to tag mobile objects such as digital cameras and audio recorders. The camera can then mark each photograph not only with a time stamp but also with the name of the photographer, since location information from the bats of the photographer and the camera tells the system who was holding that camera when the picture was taken.

All these applications give us a glimpse of how the availability of high resolution location information might redefine the ways in which we interact with the ubiquitous computing systems of the future.

2.5.4 TRIP

The problem with active tags like the badge and the bat is that they contain powered electronics. As a consequence they have a non-zero manufacturing cost and they need a battery replacement every now and then. Both of these are bad news when you want to tag gazillions of objects (e.g. every book in your home). So we devised an alternative system where the tags would cost virtually nothing to make and require no external power.

The TRIP project, or Target Recognition via Image Processing, originally developed by Henty and Jones at ORL and later reimplemented and extended by López de Ipiña [173, 172] at the University of Cambridge, makes use of special targets designed for easy optical recognition. The two inner concentric rings, which are solid, appear as ellipses regardless of the relative orientation of target and camera, and this invariant helps in locating the target and determining its size. The outer two rings, which are segmented, define by their pattern a unique identifier for each target.

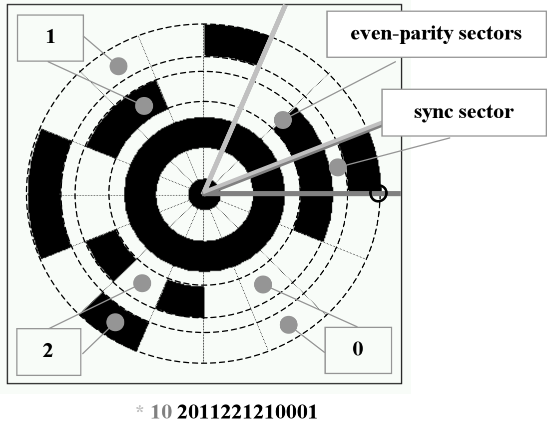

Figure 2.14. A TRIP target, showing (anticlockwise from East) the sync sector, the two redundancy sectors and the remaining 13 data sectors, for a total of 3^{13} valid codes (courtesy of AT&T Laboratories Cambridge and Diego López de Ipiña).

TRIP targets can be generated at essentially no cost on a conventional laser or inkjet printer and pasted on the objects to be tagged. Their physical size is not fixed: depending on the application they usually range in diameter from one to ten centimetres, but there is no reason why one could not draw targets over a metre wide (and I'll soon describe an application that would benefit from this). The only system constraints are on the minimum size *in pixels* of the acquired image (about 40×40), so it is acceptable to make the targets smaller if one is prepared to move the camera closer to the target or to acquire images at a higher resolution.

The encoding method used, in which black-black is reserved for the sync sector, allows for 3^{15} ($\approx$ 14 million) different codes. If desired, some of the bits of the identifier space may be used as redundancy—i.e. as an error detection code. In the current implementation, a sample of which appears in figure 2.14, two sectors are used in this way, leaving 3^{13} ($\approx$ 1.6 million) valid codes. This is a useful trade-off between the number of available codes and the robustness of the recognition: with careful tuning, the rate of false positives can be reduced to arbitrarily low values. If needed, a larger code space could be obtained by adding further segmented rings—even just a third one would yield 7^{15} ($\approx$ 4.7 trillion) codes—at the cost of requiring a larger image of the target (in pixels) for the recognition.

There are at least two ways in which such a system may be used. In one, TRIP is used as a location system. With suitable inverse geometrical transformations, the position and size of any recognized tags in the scene are transformed into geometrical coordinates with respect to the position of the camera. Depth information (i.e. distance between target and camera) can be extracted from the size of the target in the acquired image if the size of the original is known. This can be done by using some bits of the code (4 ternary sectors in the current implementation) to encode the physical radius of the printed tag, at the expense of a further reduction in the cardinality of the code space ($3^9 \approx 20,000$ codes). For absolute positioning, the camera is static[19] and its position and orientation have been accurately surveyed. Depth information could also be extracted through stereoscopic operation, by correlating the views from different cameras pointing at the same scene. Multiple cameras are desirable anyway, in order to guarantee coverage of the designated area regardless of the orientation of the tag. Such a system may be used for similar applications to the ones made possible by the Active Bat we described in section 2.5.3, such as the follow-me videophone.

In the second and simpler mode of operation, TRIP is used to detect not the precise 3D location of an object but merely the presence or absence of specific tags in the current scene. A virtual jukebox, for example, allows you to select songs by showing the appropriately tagged CDs to the camera.

An application that better illustrates the power of the system and the usefulness

[19] One could imagine relaxing this condition if the position and orientation of the moving camera were dynamically assessed using some other high precision location system, but the errors would compose and it would be difficult to produce a robust and accurate system.

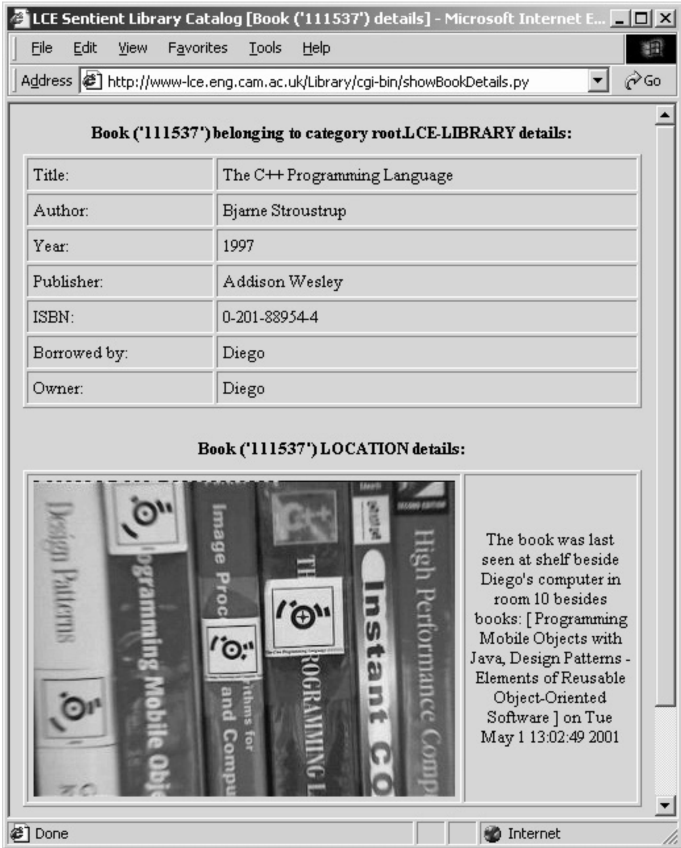

Figure 2.15. The web interface to the TRIP Sentient Library. The query shows the result of a search for Stroustrup's *The C++ Programming Language*. The bookshelf where the book was last seen is identified, and the titles of nearby books that were recognized are also mentioned. Note, in the captured video frame, a small cross in the centre of the TRIP target for the relevant book (courtesy of Diego López de Ipiña).

of being able to tag a large number of objects at no cost is the TRIP Sentient Library (figure 2.15): having tagged all the books in the lab with a TRIP target on the spine, you can walk around with a camcorder every week and film all the bookshelves. Then the videotape is postprocessed and, for each book, the system stores in a database the most recent frame containing that book. This way, when you query the database for a book, you are shown a picture of where it was last seen.

We have also thought of an application that would make use of very large targets: the Active Car Park [241]. As a matter of fact we *already* have a pretty "active" car park, even without any ubicomp additions, in so far as demand exceeds supply; it is common for staff members driving into work to try the front of the building, then the back, only to be forced to look for alternatives elsewhere—a process that the heavy morning traffic makes slow and frustrating. If we painted a TRIP target in each car bay, making them large enough so that cameras pointing at the car park from the top floor of our building could recognize them, then the system could tell which spaces were still free (because their targets would not be obstructed by a car) and this information could be consulted by users via cellular phone as they drove towards the office.

2.5.5 PEN

While the systems we have so far described have all been, one way or another, sensor systems (with an associated software back-end, sometimes of remarkable complexity), another essential component of ubiquitous computing is a facility for short range wireless communication.

Around 1994–95 we realized that the then-current infrared communication facility, included in most new laptops[20] in the form of an IrDA port [140], was inadequate for do-nothing ubicomp because of its propagation characteristics. The infrared signal follows the laws of optics and therefore it is often necessary to point the transmitter at the receiver in order to ensure a good connection. Even if reflections on walls and furniture may sometimes help propagate the signal and therefore relax the line of sight requirement, it is certainly not possible for a device on your desk to communicate with one inside your closed handbag or briefcase. We decided that radio was a much better medium to use and set out to build a system to meet the requirements of ubicomp communication.

Our aim was to build a system that could be embedded in everything as a universal communication facility. We would probably not have started this project had Bluetooth been there at the time—we would have just *used* Bluetooth ourselves, concentrating instead on ubicomp applications. But since we had to build the infrastructure ourselves, we ended up with something slightly different. While Bluetooth supports continuous media (e.g. voice) and relatively high bit rates (of the order of Mb/s), we placed a deliberate emphasis on low power consumption (for a battery lifetime of at least a few months) and consequently on short range (a few metres) and low bit rate (in the tens of kb/s).

Short range and low bit rate follow naturally from the quest for power saving, but they also have intrinsic merits. In ubicomp communications, many devices need to talk to each other simultaneously. A given frequency can only be reused when the transmitting device is out of range of whatever other device is currently using that frequency, and in this context a short range is an asset. Short range also means

[20] But seldom used by applications.

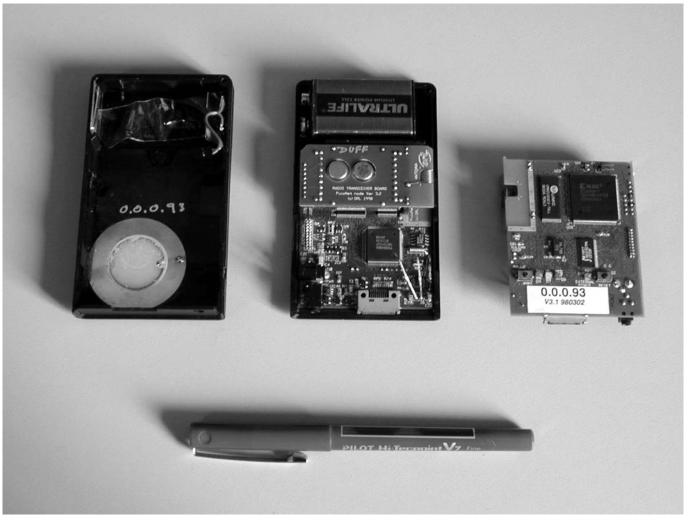

Figure 2.16. A PEN node, revealing its insides, and a lowercase pen for size comparison (courtesy of AT&T Laboratories Cambridge).

that, except for routing[21], "connectivity suggests proximity"—a useful location hint for ubicomp applications: devices that I can contact by radio are devices that my owner can see and reach. The low bit rate helps stake out an application space in which devices exchange brief nuggets of information as opposed to large lumps of data or continuous media streams. This, in turn, encourages a "mostly off" duty cycle[22] that is advantageous not only to battery lifetime but also to temporal reuse of the radio frequency, again helping towards the goal of allowing many pairs of devices to use the channel concurrently.

Our system, originally called Piconet and later renamed PEN (Prototype Embedded Network) to avoid confusion with the Bluetooth usage of the term, was first built and described by Bennett *et al.* [31]. A more recent and higher level overview of the system and its applications is offered by Jones and Hopper [146], while some more technical papers by Girling *et al.* [114, 115] describe in further detail the special low power communication protocols of the system (how do you contact a device that spends over 99.9% of its time asleep?).

[21] Not implemented in our system, although Stefanova [242] experimented with it.

[22] This may look like a paradox—it would be logical to argue that a lower bit rate forces one to occupy the channel for a longer time in order to transmit the same amount of data. However, this fails to take into account that resource availability contributes to the definition of the application space: an environment in which the channel has a capacity of megabits per second attracts different usage patterns from a 10 kb/s one. Witness the change in attitude between users of dial-up modems, ISDN, ADSL and always-on LAN-based connections.

PEN is built around a self-contained unit, the “node” (figure 2.16). In its current prototype form the node is about the size of a deck of playing cards; but in due course it ought to fit in the corner of a chip. It contains a CPU, some RAM, a generous amount of flash memory for OS and application, a radio transceiver module and some I/O ports to which one can attach peripherals such as sensors and actuators. But of course this is only the engineer’s view, which needs to be turned on its head in order to describe the real goal of this research. Nobody is interested in PEN nodes with peripherals attached to them: as per Weiser’s “disappearing computing” paradigm, what we want is ordinary objects endowed with PEN communication facilities so that they can talk to each other. The new name of our system, with its accent on *embedding* the network, is appropriate and revealing.

Figure 2.17. Cyberspice. The music starts playing as soon as you open the CD case (courtesy of AT&T Laboratories Cambridge).

Endowing everyday objects with communication capabilities can be another way to move the user interface from computers to the real world. In a popular application of PEN, called Cyberspice after the spice jar of section 2.3.1 “and a pop group that were popular at the time” [146], opening a compact disc case causes the room to be filled with hi-fi music from the corresponding album (figure 2.17). We filled the unused space under the disc holder of the standard CD case with a miniature circuit board implementing a transmit-only subset of the standard Piconet functionality. We also added a magnetic switch to detect whether the case is open or closed. When the case is opened, a message is sent from the CD case to the jukebox server, which starts playing the appropriate disc based on the identifier of that particular CD case. A flat button inside the case can be used to skip to the following track. This application demonstrates that even a very low bandwidth communication facility may be sufficient to enable useful synergies between different components of the ubicomp system.

Embedded wireless connectivity is extremely convenient for deploying an ad hoc array of sensors. As an illustration we have conducted a few practical experiments with PEN-enabled thermometers, using them to diagnose temperature problems in the home or office. The sensors can be distributed in the environment in the exact spots where one wishes to monitor the temperature—a process that would be substantially more cumbersome with wired devices. Various types of nodes from a general purpose modular toolkit cooperate to collect the temperature readings, in an excellent demonstration of ubicomp communications [146].

- The **sensor** nodes are battery-powered. Once deployed they spend most of their time asleep, but wake up regularly to sample the temperature and talk to other nodes.

- The **cache** nodes are mains-powered, so they can afford to be listening all the time. Whenever a sensor produces a reading, the cache node stores it, so that it can be retrieved later by another node querying the temperature even while the sensor is asleep. A cache node may serve several sensor nodes, and it is desirable that each sensor node be covered by at least one cache node.

- The **logger** node collects all the new readings it finds from sensor or cache nodes and stores them in its long-term memory. It keeps a log of all past readings as opposed to just the most recent ones. It is battery-powered so that it can be carried around to collect data automatically from many static nodes as the user walks past them.

- The **download** node is interfaced to a PC. It sucks data from other nodes and transfers it to the PC for processing (statistical analysis, graph plotting etc.).

We have used this type of setup to monitor the efficiency of the air conditioning system in our machine room when it appeared that some of the computers were not being adequately cooled. Compared to the paper chart recorder depicted in figure 2.18 the PEN sensors are smaller, so they interfere much less with the chilled air flow they are supposed to measure; they are cheaper, so more of them can be deployed; and they produce a digital output that is much more convenient to postprocess. They can for example be used to produce a single table or graph correlating the readings of several sensors at different points in the room.

The convenience of tetherless deployment is not limited to input devices such as sensors: it is equally advantageous for output devices such as displays. We have

Figure 2.18. A paper chart temperature recorder and a PEN temperature sensor (courtesy of AT&T Laboratories Cambridge).

Figure 2.19. A dynamic room tag combining PEN, a cholesteric display and a data feed from the Active Badge system. Because they are wireless, such tags can be very easily deployed throughout a building (courtesy of AT&T Laboratories Cambridge).

built a low power wireless display by coupling a PEN node to a special cholesteric LCD that requires energy only to *flip* the state of its pixels, not to maintain it—in other words, once you have written a bitmap to the display, it will stay on indefinitely even if you remove the battery. This allows us to treat the device somewhat like a printed tag, and indeed we have used this type of PEN node as a door tag (figure 2.19). Compared to a traditional display there is still no noticeable energy penalty in changing the information on the display a few times per day, so we update the door tag dynamically with information from the Active Badge (or Bat) system, showing who is inside the room at the time. We can also display a short message from the owner of the room, such as “back later” or “please do not disturb”.

2.6 Security issues

The preceding sections of this chapter have given us a glimpse of what the ubicomp-enabled future might perhaps bring. As Weiser noted in his seminal paper, we don’t really know what’s coming [259]:

> Neither an explication of the principles of ubiquitous computing nor a list of the technologies involved really gives a sense of what it would be like to live in a world full of invisible widgets. To extrapolate from today’s rudimentary fragments of embodied virtuality resembles an attempt to predict the publication of Finnegan’s Wake after just having invented writing on clay tablets. Nevertheless the effort is probably worthwhile.

What about security? We don’t know in what ways the ubicomp scenario can be abused by ingenious attackers, and we don’t know who the attackers are going to be. You’ll find that they probably won’t be limited to the computer villains to whom the press incorrectly refers as “hackers”. Here too, therefore, the imagination effort will be worthwhile, and an important step will be to identify what exactly we want to protect. Maybe something we take for granted right now…

For example, when store loyalty cards were introduced, many people did not see any drawbacks in a system that gave them a discount in exchange for telling the merchant what they had bought. “After all, the merchant already knows what I bought and, unless I pay cash, it can even link the till receipt to my cheque or credit card account.” This is true in theory; but, in practice, it requires a little too much work to be worth doing. The store card takes away the guesswork and makes it possible for the merchant to build an accurate profile of your purchasing pattern at very little cost. A crucial consequence is that such profiles can be traded among merchants *and merged* to build disturbingly exhaustive dossiers. Imagine the dossier obtained by combining the log of all your transactions with supermarkets, airlines, bookstores, car hire outfits, banks, highways, phone companies and so on. Even a very partial subset of it will reveal where you were on any given day, whether and when your household had any guests, whether your overall diet is healthy, whether you are addicted to tobacco or alcohol[23], what kind of books, newspapers and videos interest you, what dangerous sports you practise and so

[23] Perfectly legal in many jurisdictions, but you might consider your vices to be your own business.

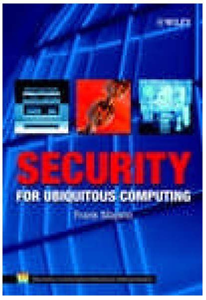

Frank Stajano (University of Cambridge)
Security for Ubiquitous Computing
John Wiley and Sons, Ltd
Wiley Series in Communications Networking & Distributed Systems
ISBN: 0-470-84493-0
Hardcover; pp. 267 (xx + 247)
Publication date: 2002-02-12
RRP: 34.95 GBP (UK); 59 EUR (rest of Europe); 60 USD (USA)

http://www-lce.eng.cam.ac.uk/~fms27/secubicomp/