WHITE PAPER SERIES / EDITION 1

AUTO-ID LABS

BUSINESS PROCESSES & APPLICATIONS
SOFTWARE & NETWORK
HARDWARE

AUTOIDLABS-WP-HARDWARE-020

An Overview of Location and Identity Technologies

Hugo Mallinson
Auto-ID Center, Institute for Manufacturing
Department of Engineering, University of Cambridge

www.autoidlabs.org

Abstract

In automation systems, the sensing of location and identity is essential to provide the interface between the physical space and the computer logic which runs the system. It is important to understand, however, that the two are not necessarily synonymous, and most technologies provide each datum to varying degrees. This paper provides a survey of some existing sensing technologies, and positions them in a framework within the general sensing space.

1. Introduction

1.1 Automated Environments

The capability to automate manufacturing has made a significant difference to manufacturers in this century; computerized systems to sort and manage goods have now become essential in many industries. Robots now perform many tasks which used to require humans. With the rise of computerization it has become increasingly important to properly identify items in production. Robots provide their controlling computers with access to the physical world, receiving feedback from a variety of electronic sensors. Having these interfaces between machine and physical realms allows computers to not only perform tasks, but to receive feedback from the environment and use it to modify their future actions. Properly programmed and equipped with adequate feedback, robot systems can run almost indefinitely. Environments which use robots and equipment which is configured to operate autonomously are known as automated environments.

Automated environments promise great benefits for companies which use them. A computer can sort and route items precisely and repeatably, far better than a human. Having computers which can perform these sorting tasks makes assembly lines more flexible because many different items can share a single assembly line, with the computer keeping track of each item and its specific needs. Computers are capable of managing huge amounts of data at very high speeds and make it possible to track all of the products in an assembly line individually. Specific actions can then be performed on any given product, even in a fast-moving environment. A system's output is only as good as the input it is given, however. Computers can keep track of where an object has been, where it needs to go and what needs to happen to it, but the initial information about which object is being handled needs to be accurate. Automated systems use various sensors to connect the physical world to the computer; these sensors must provide two pieces of data for each object in the system: its identity and its location.

1.2 Identity and Location

This paper will discuss the differences between identity and location; it is important to realize that these two are not necessarily linked. Humans can easily identify objects and judge their position visually and can therefore think of the two attributes as synonymous. While often related, the two are in fact discrete. We often see something but can't make out exactly what it is, or recognize the voice of a friend but not know where he is. It is possible to know many things about an object without knowing its location, and it is possible to know exactly where an item is without knowing its identity.

In an automated system knowledge about location could come from physical sensors like switches or from optical sensors which indicate that some object is pressing the switch or breaking the light beam. Identity could be determined by reading a serial number or by measuring some physical attribute like weight. Both location and identity can be measured to varying degrees of accuracy. Location is never exact; it is impossible to say where an object is with absolute certainty. An oft-quoted riddle asks "How long is the coastline of Britain?" The answer is infinite, because the more closely one examines the coastline (looking at smaller and smaller bays and inlets) the longer it becomes. The same logic applies to any distance measurement: it can never be known absolutely, only to the degree of accuracy to which it was measured. Identity is more discrete — an object may certainly be identified correctly — but there are levels of identity as well. A can of Coca-Cola can be identified as cylindrical, and further as red, and further still as weighing less than a kilogram, and further still as being a can of a liquid, and finally as being a specific can of Coca-Cola.

In order to be useful, in many applications these two attributes must be linked, and most systems rely to some extent on inferring one from the other. For example, if an object's location and identity are known and its location changes, it is sometimes safe to assume that the object at the new location has the same identity as the old object.

2. Location Sensing

There are generally two ways to locate objects: by their inherent physical properties or by a tag which has been specifically installed in order to track them. This section examines methods of both types for determining item locations. It introduces a classification method for these technologies which describes them in terms of three attributes, and plots them on three corresponding axes. This arrangement makes an immediate visual connection to the technologies and helps explain where each fits in relation to the others. These techniques are examined to determine their fitness to be integrated into an automated system/environment to enhance the acquisition of location data.

2.1 Tagging

A tag is a distinguishing mark added to an item to make it machine-readable. It could be a printed barcode, an electronic tag, or even a mark with a special pen. A tagging system consists of tags, which are attached to the items being tracked, and some sort of electronic reader which interacts with the tag to produce ist identifying number. Electronic tagging works in the same way as handing out name tags at a meeting, by providing a standard way for people's names to be known. Anyone who wants to determine a person's name simply looks on that person's lapel and reads the tag; the meeting conforms to a prior protocol known by both the reader and the tag (in this case that the string of text on the tag is the person's name).

2.2 A Method of Classifying Location Systems

To give order and provide a basis for comparison of location technologies, a classification system was devised. The technologies can be categorized by three independent attributes: the physical medium through which they work, the characteristics of the medium which they measure, and the manner in which those characteristics are processed.

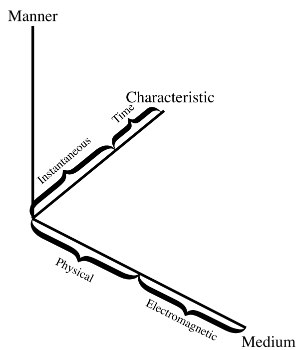

Medium The medium through which the data are collected (and therefore in which the object's position is measured). The two media of practical use are the Physical, which carries soundwaves and tactile sensation, and the Electromagnetic, which carries radio and light waves.

Characteristic There are two main characteristics measured: instantaneous intensity and the time delay between two events. Intensity can be gradual, as with a scale, or binary, as with a simple switch.

Manner This axis can also be thought of as the complexity of the measurement. In some cases one simply needs binary data, e.g. "Is the switch on or off?" In other situations more subtle gradations are necessary: "How hard is the scale being pressed?" and still other situations call for more complex arrays of data: "How hard is each of these *n* scales being pressed?"

2.3 Tactile Sensing

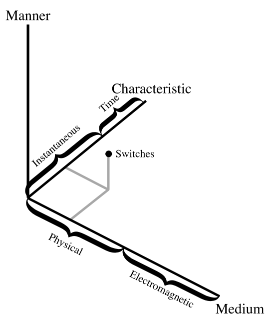

The simplest form of location sensing is tactile, provided by switches. A tactile sensor works in the physical medium and measures the (binary) intensity of pressure on a point or region. If a switch is placed in the middle of a conveyor belt, when it is triggered the system knows that an object is at that point.

Switches fall very near the origin of the three-axis model: physical medium, low-complexity characteristics, low-complexity processing. From this point there are three ways to enhance this sensor. The complexity of the processing could be changed, for example by adding a second switch a small distance from the first.

By measuring the interval between signals from the two switches the velocity of the object could be determined. The complexity of the measurement could also be enhanced by replacing the switch with a scale. The data would now be continuously varying over some range and could be used to sort items by weight. It should be noted that any switch provides weight information, since it requires some activation force, but its output is still binary, i.e. the item either is or is not heavy enough to activate the switch.

A third way to alter a switch would be to move to the electromagnetic medium and use a light- or radio-wave detector.

2.3.1 Proximity Sensing

When an item is detected by a sensor, whether a microphone, a break-beam, a switch, or other device, the item's location is immediately known to within the range of the sensor. In the case of a break-beam some part of the item must lie somewhere along the beam, for the microphone the object must be making noise within range of the transducer. This is the most elementary form of localization in that it requires the least calculation — an object is detected and its location is immediately known to within a known error. There are gradations of this accuracy: a break-beam can be sensitive to a fraction of a centimeter, whereas the microphone does not have nearly as sharp an edge to its reading field; it can reliably locate items within its range, i.e. a noise very nearby, but those at the furthest limit will not be detected reliably.

The ActiveBadge

The ActiveBadge is a good example of a proximity system. Developed at Olivetti Research Ltd., UK, it was designed to locate and track people and objects within an office building. The system consists of a network of infrared receivers installed in rooms throughout the tracking area and portable badges small enough (roughly 55x55x7mm and 40g) to be worn by a user. The badge broadcasts a unique 48-bit string every 15 seconds via an infrared transmitter.

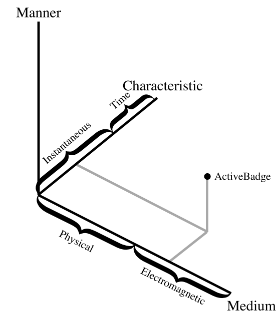

The infrared signal reflects off various surfaces in the room with sufficient intensity that a wellplaced receiver can receive a transmission from almost anywhere in the room. The walls of the room act as a natural barrier for the signal, so that stray signals are very unlikely. Thus the system can locate tags to within a single confined space. Since the badges transmit a signal they require batteries, whose lifetime varies with transmission frequency [15].

2.4 Signal Strength

Any transmitted signal, whether by sound, radio, or light, will degrade as it travels further from its source. If the medium through which the signal travels is relatively uniform this degradation can be modeled and, given appropriate information about the configuration of the transmitter and receiver, the distance between the two can be estimated. In practice this technique is difficult to implement, especially indoors where walls, tables, electronic devices and other objects interfere with transmissions and make their degradation very hard to predict. A system based on signal-strength measurement would fall farther along the characteristic axis than the proximity system because it has to establish a scalar quantity rather than a binary one.

Signal Strength with Added Processing

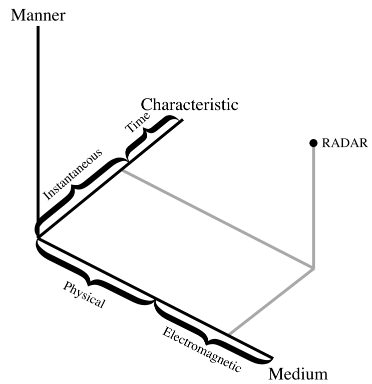

The RADAR system, developed by Microsoft Research, embraces this deviation of signal strength from the theoretical rather than trying to minimize it. The system consists of a number of 802.11b wireless ethernet base units arrayed around the space being instrumented. The item being tracked, in this case a laptop with an 802.11b card, is moved around the space and samples of the signal strength from each station are taken at set intervals. Once these data are recorded, the problem of locating a receiver given its signal data becomes a matter of statistics: interpolating the most likely location for the receiver from the existing data. This technique increases accuracy over the theoretical model since it accounts for any variations in the environment from the theoretical ideal. In the test environment used by Bahl et al the accuracy increased with the addition of the empirical data processing. The error for the 50th percentile of the measured data went from 4.3m for the theoretically calculated location to 2.94m for the empirically calculated location. This is a signal strength system extended along the Manner axis [3].

Visual Acquisition

Video cameras are an array of photo-sensitive cells which measure the intensity of electromagnetic radiation in the visible light spectrum. To a computer, then, the problem of video tracking becomes one of handling an array of variable inputs and performing additional post-processing. Items can be tracked visually, by a video camera installed at the proper vantage point. Video surveillance requires little infrastructure, but puts great demands on the processing hardware which must locate the items to be tracked within the image. In environments where minimal impact is desired, video can be an essential asset. Trakus, Inc. uses a combination of inertial sensors and video to collect data for live television and Internet broadcasts of ice hockey games [14]. Many tagging applications, however, lack the consistent clear line of sight between camera and object which video requires to be of much use. And video does not scale naturally, as it must be carefully calibrated for its environment and the cameras cannot simply be moved or zoomed in or out without extensive recalibration.

2.5 Time Measurements

A distance can also be calculated by measuring how fast it is covered by an object of known velocity.

Time-of-Flight

This method measures distance by measuring the time it takes a pulse to travel between the two objects being measured. The pulses are usually ultrasonic or Radio Frequency (RF). A pulse is transmitted to a receiver, which returns the pulse to the transmitter. Calculating the distances requires knowing the exact time of flight. This time can be found by:

- synchronizing the timing of the transmitting and receiving nodes and finding the difference between the transmit and receive times or
- timing the entire round-trip signal travel time, subtracting the receiver's processing time, and dividing in half.

If pulses are sent from a series of base stations to the object being located and their roundtrip times are calculated, then the object's position can be determined.

Trilateration

Determining location in a multi-dimensional environment requires more than a single distance measurement. Readers are probably familiar with the concept of triangulation: given the bearing of an object from two known positions, the object's two-dimensional position can be determined. Trilateration is a similar process but with distances. If the distances to an object from three known (non-colinear) locations can be determined, then the object's location is known in two dimensions. If three-dimensional location is required, a fourth measurement must be made. It is sometimes possible to make do with fewer measurements. two distances will locate an object in two dimensions at one of two positions, and if one of them is known to be impossible (in a place where the object is known not to be) then it can be disregarded. Given that an object can be located based on its distance from known points, let us examine methods of determining those distances.

Time Difference of Arrival

Time Difference of Arrival (TDOA) measures Time-of-Flight (ToF) in a one-tomany configuration. The object sends out a pulse of ultrasound and the base stations, which are synchronized with each other, record the time each receives it. Since the base stations cannot know when the pulse was sent, only the relative reception times, they must perform additional calculations to determine the transmitter's location. The location can still be determined with the same number of stations as explained in the previous section.

2.5.1 RF Time-of-Flight

PinPoint 3D-iD

The PinPoint 3D-iD system uses time of flight calculations to locate its nodes, but instead of using ultrasound pulses, it uses a transflected RF pulse.

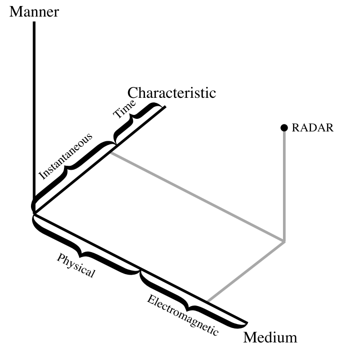

The system uses radio tags and a series of base stations, which can each be attached to up to 16 antennas, which in turn send out RF pulses to trigger the tags. The base station sends out a pulse on the 2.4GHz frequency band, which the tag receives and modulates with its unique ID and sends back to the base station. The base station subtracts the known processing delays in the tag and its own hardware to get the total travel time, halves it, and multiplies by the speed of light (the propagation time of an electromagnetic wave) to get the distance to the tag.

2.5.1.1 Ultra-Wideband

Ultra-wideband (UWB) technology can be used in RF communications to decrease interference, lower transmission energy and, in localization scenarios, to increase positioning accuracy. UWB radios do not use sinusoidal oscillators like conventional transmitters, rather they transmit data in very short bursts which, by virtue of their extremely short transmission time, have very wide bandwidth. Signals with wide bandwidth share all the advantages of spread-spectrum for noise immunity and minimizing interference: by using many frequency bands the chances of encountering constant interference are lessened. Unlike spreadspectrum however, since all of the broadcast power is transmitted simultaneously the power at any given frequency is extremely small, so a comparatively large amount of power can be put into a signal without affecting other devices broadcasting within range of the UWB device. The pulses transmitted by UWB equipment are extremely sharp. A well-defined peak is easy for the receiver to locate, and therefore the travel time of UWB pulses can be measured to much higher accuracy than conventional RF pulses. Current systems can reach 1cm accuracy over a 50-60m range [1].

Ætherwire Inc.

Ætherwire Inc. designs a line of UWB products very similar in infrastructure to the MIT Cricket explained later. These "localizers" use time-of-flight calculations to find their range to all their surrounding localizers. Thus a distributed network is created where each node knows its location in terms of its surrounding peers [1].

2.5.2 Time in the Physical Medium

Distances can also be measured by calculating the propagation time of sound waves through the air. This is generally done with ultrasonic transducers.

The ActiveBat

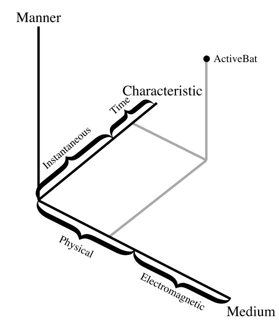

The ActiveBat was developed by Olivetti Research to provide greater accuracy than the ActiveBadge. The transmitter is smaller than the ActiveBadge, and the protocol is more complex. It consists of a wearable transceiver (5x3x2cm, 35g) and a network of nodes designed to be installed above a dropped ceiling. A base station sends an RF interrogation to each Bat in turn, asking it to identify itself, and simultaneously sends a signal over the wired network of nodes to synchronize their clocks. When a Bat receives its interrogation it responds with an unencoded ultrasonic pulse. The nodes receive the pulse at varying times and their distance to the Bat is calculated based on the speed of sound in air. The system waits a preset interval to allow the ultrasonic echos to die out before interrogating the next Bat [4].

The Bat system includes processes to automatically register new Bats which enter its space. It can also use intelligent polling schemes — increasing polling frequency for Bats which are likely to move often and decreasing it for those likely to remain stationary, e.g. Bats located at workstations. The protocol also allows for radio transmissions from the Bat to the base station — an action button on the Bat can provide feedback to the system. Since most solid objects provide a good degree of sound deadening, a Bat placed on the side of an object will only be heard by receivers not in the shadow of the object. These receivers can therefore obtain a rough approximation of the object's orientation. With 100 receivers covering an area of $280m^2$, 90% of readings can determine the orientation to within 60°. Over this area 95% of readings are accurate to within 9cm [4].

The Cricket

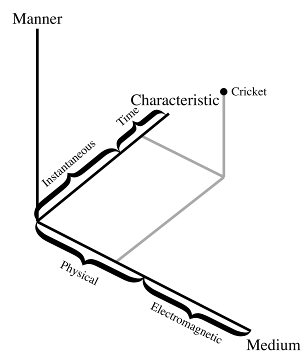

The Cricket system developed by the Networks and Mobile Systems group at MIT's Lab for Computer Science uses the same RF/ultrasound technique as the AT&T Bat, but with a decentralized topology. The Cricket system has no central processor which locates its nodes, rather, each node locates itself in terms of its closest nodes. The system consists of beacons, which identify a space, and listeners, which are attached to objects that need to locate themselves in the space. The beacons transmit a simultaneous RF and ultrasound pulse at random intervals to minimize interference with other beacons. The listeners are triggered by the RF pulse, when they begin waiting for the ultrasound (travelling at the speed of light) to appear. Once the ultrasonic burst (at the speed of sound) arrives, the listeners can calculate the distance to the broadcasting beacon. After finding a few beacons the Cricket knows which is closest and therefore which space it is in [10].

Properly the Cricket technology is a Proximity localization technique, but its technology could be easily used for trilateration and so it is included here.

2.6 Higher-Level Processing

2.6.1 Dead Reckoning

In a production environment there are myriad possible snags in locating items. Sources of interference can change and move and the nature of radio waves means that there can be "null points" caused by reflecting waves canceling each other out where an item cannot be located. Given these transient errors there can be a great advantage to using some form of dead reckoning in the software which handles locating the items. Even simple logic along the lines of "This tag has been moving with velocity $\vec{v}$ through 40 updates and now it's missed an update, I'll assume it's still moving at $\vec{v}$" can prove useful in many environments.

The PinPoint system uses dead reckoning of this sort in some of its systems. By making educated guesses about tags based on knowledge of the environment, it is able to make the most of partial reads, where a tag's range from only one or two of the necessary readers can be calculated, and situations where a tag cannot be read at all. It can, for example, say that a missing tag was last seen outside a door, and so it is very likely inside the room even though there are no readers inside to verify this [16].

2.6.2 Topological Awareness

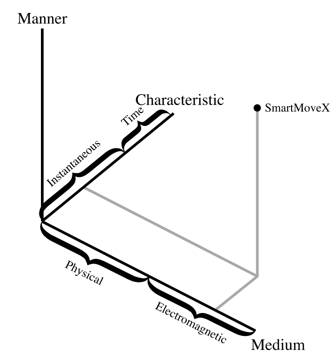

The SmartMoveX system designed at Microsoft works in a similar fashion to RADAR, but substitutes small radio transmitters for 802.11b hardware. Its major innovation, however, is its intelligent mapping of the space (A further extension along the Manner axis). Along with the node data used in the RADAR system, it stores a map which represents the physical connection of the nodes, showing for example which nodes are in joined rooms, and whichhave walls between them. This provides an additional level of sanity-checking to results, as it will catch objects that appear to move through walls or jump between nodes not immediately connected [8].

2.6.3 Inertial Tracking

For situations where a small footprint is not essential, extremely high precision can be achieved with the additional use of inertial tracking modules. For example, an accelerometer mounted inside an object could have its output integrated to determine the object's current speed, and integrated again to find its distance traveled. Intersense Technologies makes an inertially-tracked 6 degree-of-freedom input device. Ultrasonic time-of-flight calculations are used to give initial position data and to keep the inertial sensors from drifting. The system, using a grid of 4 sensors in a 2.5m square array, can cover an area 1.3m below the sensors and give 1.5mm accuracy. Presumably the accuracy would fall off roughly linearly with increased distance from the sensors [7].

Inertial tracking is currently quite expensive, and would only be suitable for applications where extremely high accuracy was essential. It is popular in motion capture applications for film and virtual environments, where small gestures need to be accurately recorded. The onboard computation is intensive and would require a large power supply and careful construction and mounting, as compared to an active radio tag or similar technology [7].

3. Identity Sensing

This section examines two leading technologies for providing identity data: barcodes and RFID. Combined with location data provided by the technologies in the previous section these can be used to track items in an automated environment.

3.1 Identity Sensing with Tags

As with location, identity sensing can be performed with or without tags. In a situation where tags are not used, items can be identified by their physical attributes. These measurements can be made by active sensors like range finders which can measure the size of an item, or by more sophisticated electronics like video cameras which can be used to identify an object by color, shape, size, or even by recognizing individual features like a person's face. Physical constraints can also be used to identify items by process of elimination; for example, letting only objects of a certain size pass along a conveyor belt for.

These sensing techniques are generally not flexible enough for use in automated environments, and the remainder of this discussion will focus tagging. Two tagging technologies are popularly used in industrial applications: barcodes and RFID [9].

3.1.1 Barcodes

Physically, a barcode is most often made up of a group of parallel black lines on a white field. Variations in the width or length of the lines are use to encode data. They are commonly seen in the Universal Product Code (UPC) labels on packages in stores and on letters (they are used to sort letters in the UK, USA, and France, among others[13]), as well as price tags, library books, identity cards and many other items which must be machine-readable. The code is read by shining high-intensity light onto the barcode and then detecting the reflected areas of dark and light. Two-dimensional, or "matrix" barcodes also exist which use a grid of points instead of an array of lines. These can store more data in a given amount of space, but require much more elaborate decoding hardware than the traditional code. The barcode as it is seen today has been in use since the mid-1960s. Originally it was used to track railroad cars; it was then introduced into factories to track finished goods [11]. Barcodes are flat and opaque. The range at which they can be read varies with the size of the barcode and reader used. Readers for standard UPC-sized labels as seen in supermarkets can work properly as far away as a meter from the barcode [12]. Some readers also use small cameras which take an image of the barcode and process it in software, so their range is potentially limited only by the lens used with the camera and the resolution of its CCD[1]. However, as range increases, field of view decreases for a given lens, which means that the barcode must be placed ever more precisely in front of the reader. Because they can be printed in standard inks on almost any material, barcodes are extremely cheap; if an item is already tagged with an inventory sticker, the costs of printing a machine-readable barcode along with the human-readable text are no more than that of adding plain text to the same label. Barcode readers can read the codes in most orientations, but they need a clear line of sight and their angle to the surface cannot be too oblique.

3.1.2 RFID

Increasingly, automated environments are using RFID systems to provide location and identity information [9]. An RFID system consists of tags, which are generally small and attached to something that needs to be tracked, and readers, which interrogate the tags and read their data.

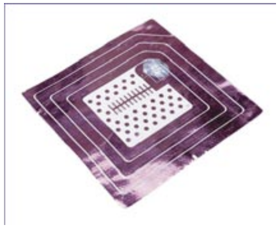

Fig. 1: An RFID tag

An RFID tag (figure 1) consists of an antenna and a microchip mounted together in a substrate which is attached to the item being tracked.

[1] Charge Coupled Device, the module which converts light into electrical signals in a video camera.

The RFID reader (figure 2) transmits an electromagnetic wave which powers the tag, allowing it to transmit back to the reader. From the received signal the reader determines the tag's ID. This process happens in a fraction of a second, the time depending on the powering up of the tag, the presence of radio interference between tag and reader, the processing of the signal in the reader to determine the transmitted data, and factors such as obstructions or the number of tags being read simultaneously. Tags can be made in a range of sizes and materials to be attached to any number of objects. Their range tends to decrease with the size of their antenna, but if short range is acceptable, tags are currently available in a form factor 56 x 4.75mm with a range of 25cm[5]. RFID systems are primarily used in one of four frequency ranges, as explained in table 1.

Fig. 2: An RFID reader

A casual user of an RFID system will note that a tag placed in front of the reader is read and its identity revealed. It is therefore commonly assumed that RFID systems provide simultaneous location and identity information. This is not always the case — neither the precision nor the accuracy of an RFID read is necessarily predictable. RFID systems can show both false positive[3] and false negative reads even when tags are very near to the reader. Furthermore, the resolution of the location reading is related to the range of the combination of reader and tag, which is in turn dependent on a host of factors including, for example, the size of the reader and tag's antennas, their design, the power output of the reader and tag, and the presence of noise or interference in the environment. Additionally, readers do not have a facility for saying where within their range a tag exists. The antennas simply receive signals, without knowing where they come from. It is quite possible for a tag outside the "range" of a reader to receive power from another reader and then broadcast its identity in the direction of the first reader, resulting in a read. What can seem at first glance to be a simple binary sensor is in fact a very complex system where subtle changes in environment and methodology can have a profound impact on behavior. It is for these reasons that RFID often needs to be augmented to provide adequate location data.

[2] The first "agile" readers, as multi-frequency readers are known, have been developed by ThingMagic (http://www.thingmagic.com).
They are currently being manufactured by Tyco under the Sensormatic brand (http://www.sensormatic.com).

[3] False positive reads are of two types: detecting a tag when none is present, and detecting a tag but misreading its ID sequence. The transmissions from a tag are unlikely to be simulated by random noise, making the first sort of error unlikely. Most tags employ a Cyclic Redundancy Check (CRC) code to confirm proper transmission of their ID, making the second type of error also very unlikely. It is possible for the CRC bits to be misread in such a way as to confirm an incorrect transmission, but this is even less likely. There is a caveat to the first type: detecting a tag where one is but which is outside the "normal" range of the reader. It is possible for a reader to receive a transmission from a tag which it could not have powered itself or whose transmission it could not normally have received. The tag could have received power from another reader, or its transmission could have bounced off some obstruction in the environment. In either situation the reader will show a read, but the tag will not have been in the area normally read by the reader and any assumptions about the tag's position made in software may then be in error.

RFID technology is increasing popular in a variety of applications. The Auto-ID Center[2] is developing an open protocol for item tracking which builds on RFID with other technologies and is being targeted to RFID environments in a wide range of industries. One of its goals is to introduce economies of scale into RFID tag production to make possible the production of cheaper tags which will enable the adoption of RFID technology by companies with even the largest quantities of relatively low-value taggable assets. With this potential increase in RFID use it is important to understand the limitations as well as the benefits of the technology more completely.

Frequency Designation	Frequency Band	Description
LF	125kHz	Low frequency signals require long antennas and are generally larger and more expensive than higher-frequency alternatives. These tags are inductively powered – their power falls off as $1/d^3$ – the reader power must increase greatly for modest improvements in range. LF signals are less prone to interference especially from biological tissue and liquids, so the tags are commonly used in applications like animal tracking. The frequency band is available in most countries.
HF	13.56MHz	Tags which operate at this frequency can have shorter antennas than LF tags and greater range. This frequency band is also available worldwide, so the tags can be used legally in any country. This is important for manufacturers who work in many coutries and want to standardize their equipment or for goods which are tagged (as in section 3) and then shipped internationally. These tags are also inductively powered and are limitedby the same physics as the 125kHz tags.
UHF	915MHz (US) 868MHz (Europe, Japan, Australia)	A higher frequency than 13.56 means that these tags are potentially more powerful and can therefore have a greater range. With higher frequency comes shorter wavelength (around 33cm), however, and this makes them more susceptible to interference. Two different bands are available in the US and Europe (although the proximity of the frequencies makes the tags' physical characteristics practically identical), which makes it impossible to use a single tag in an application which must work in both regions. While it is possible to manufacture readers which read both tags, they are still in the early stages of development.[2]
UHF	2,45GHz	The top of the RFID spectrum is also available in pretty much every country. These tags have an even shorter wavelength, only about 12cm, making

		potential interference even more likely than with the other UHF tags. This frequency band is shared by other technologies, including Bluetooth and 802.11b as well as many other short-range radio devices.

Table 1: RFID tag comparison

3.1.3 A Tagging Example

To demonstrate the use of tags in an automated environment, let us examine two scenarios for a toy factory. The two factories will be called A and B, where the first does not use tags and the second does. The factories paint and package three kinds of toys: blue rattles wrapped in clear plastic, red rattles also packaged in clear plastic, and red lobsters, packed in little cardboard boxes.

Initially both factories must identify which are which and sort them accordingly. This is done with physical sensors. Plant A then allocates some rattles to be red and puts them on one assembly line to be painted, it puts the remaining rattles on a second assembly line to be painted blue, and the lobsters on a third assembly line. Each assembly line paints and then packages its product.

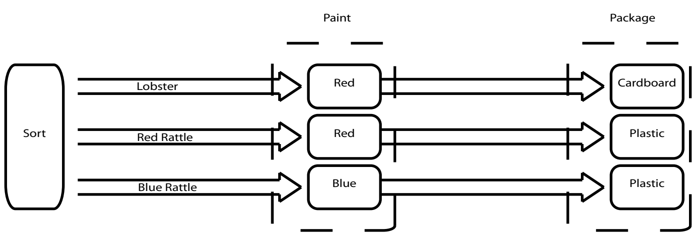

Fig. 3: Factory A

In Plant B, each toy is given a tag which identifies it: a small radio chip glued to the underside. The four processing steps (two colors of paint and two packages) are shared among the three toys and they can all be put on a single conveyor, since the radio tags allow the machines to know immediately what toy each is.

As needed, objects are routed to different conveyors. First the lobsters and the rattles which are to be painted red are send to the red painting robot and the remaining rattles are sent to the blue painting robot. A tag reader reads each item coming out of the painting robots and routes it to the correct packaging system, either plastic or cardboard. In a manufacturing environment, tagging is even easier because tags can be installed as part of the production process, so that step would be eliminated from this example. The tags give each toy a uniform identifier which the system can use to track it. If an order comes in for blue lobsters, the system simply has to be told to route some of the lobsters to the blue machine instead of the red. Likewise, if a promotional tie-in requires rattles in cardboard boxes the system can be told to send the rattles to the cardboard packaging system.

What tagging brings to the manufacturing process is flexibility. In this situation plant A would have to add a series of physical sensors after the painting step to identify each item again and route it to the correct packaging station. The use of radio tags means that once the toys' identities are known they are not lost throughout the manufacturing process. Once items are tagged the problems of tracking them can be faced.

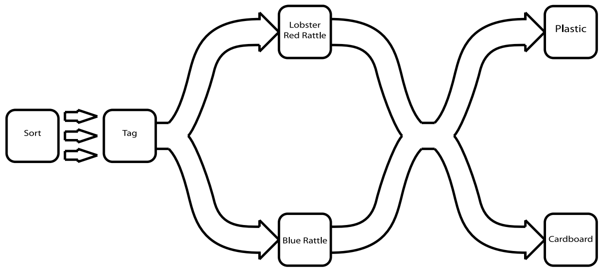

Figure 4: Factory B

3.2 Tracking

Tracking means keeping a record of where a particular object was at a particular time. To do this a system needs to know the object's identity and its location, which come from tagging. Uniform tags mean a single reader will read all items, so tracking requires only deploying one type of sensor. This sensor will give the identity of the object and its location so far as the object is within the range of the sensor. These ranges can vary greatly, however. Tracking is the next organizational step above tagging; it takes the information given by tags and organizes it in a meaningful fashion.

It can sometimes be helpful for such a system to make assumptions about objects rather than measuring each one in order to reduce the processing requirements or sensor cost for a system. If a series of sensors has identified an item on a conveyor belt as type X, it could be tracked subsequently simply by knowing the time and location at which it was detected and the speed of the conveyor belt. Once items have been identified in this way they can be sorted, for example by switching them to a different conveyor belt on an assembly line.

3.2.1 Tracking with RFID and Barcodes

Used in a tracking application RFID has some attributes which should be kept in mind: since RFID tags communicate with their readers via electromagnetic waves rather than visible light they can do so through more materials than barcodes. This means that the orientation of tagged items with respect to the reader is less important and the range of the reader is generally greater. Also, because of its use of radio waves as its medium, rather than visible light, RFID problems can be less intuitive to debug than those of barcode-based systems. Since a barcode reader shines a visible beam of light across the barcode, it is easy to see when it should be reading and so immediately noticeable when the reader illuminates the code correctly but cannot read it. RFID does not have this visual feedback and it is harder to know when a tag is in a position where it "should" be read.

RFID's advantages, especially its ability to operate without a line of sight connection to the tag, generally outweigh its drawbacks. Both its growing popularity and its somewhat unintuitive nature make it a very important technology to study. No other technology provides its combination of location and identity. To give some sense for the potential applications of RFID, three possible (and somewhat idealistic) applications of the technology will be discussed.

3.3 Applications

3.3.1 Factory Automation

In a soft drink factory sheets of aluminum and cartons of syrup are the process inputs. The cartons have been tagged by their manufacturer so as they come into the factory they are counted and recorded; then the system can sort them by flavor and put them in the correct tubs for mixing. The sheet aluminum is pressed into cans which are tagged. The system decides which cans will contain which drinks and as the cans move along the assembly line all of the "Diet" cans are filled with the diet drink, and all of the "Regular" cans are filled with the regular drink. Farther along they are painted according to the flavor they contain. Then they are packaged. Some stores want packages of all regular and all diet, and some want gift packs that have half and half. The robot which packs the cans runs a reader along the finished cans. When it finds the type it needs it picks it up and packs it into a crate. Because of the flexibility afforded by a tag-based system, it is equally easy to create any style of crate as any other. Each time a can is scanned by a reader the tag ID string is stored in a database along with the time and the name of the reader which read it. With these data failures in the assembly line can be spotted or a specific can located.

3.3.2 Warehouse Location Monitoring

Once product has been delivered it must be stored. Warehouses can have hundreds of shipments arriving in a day and some of them can be mislabeled or misplaced. If the contents of these shipments carry RFID tags, then the true contents can be determined. The warehouse has readers installed at high-traffic points. As the shipments are moved around the warehouse the readers will record the fact that item x passed point y at time t and store this information in a central database. The database can confirm that items have been stored in the appropriate place in the warehouse.

The readers installed around the warehouse can give a general sense of product locations; in order to determine precisely which items are where a worker can run a tag reader along a shelf and see exactly what has been stored there and compare these readings of the actual contents with the labels on the packaging. This process can be automated, say by a small robot which moves autonomously around the warehouse and scans each pallet it comes to, recording its location and contents. Scanning packages in this way closes the feedback loop for the warehouse manager, confirming where all the product actually is. When the warehouse needs to ship out a pallet of a given product, it can be found immediately.

3.3.3 Point of Sale

Since retail items come already tagged from the distributor it is a small Stepp for retailers to use these tags to track their merchandise. Because orientation is less important with tags than with barcodes, it is possible to mount readers on shelves to monitor their contents. Retail locations pose particular problems, however, because they must balance the needs of the buyer and the seller. Metal shelving is problematic for RFID since it interferes with the EM signals used by the tag and reader so a single reader cannot be used to read a large block of shelves. There are ways around this, including installing more readers or using different shelving systems, but retrofitting RFID technology into stores still requires careful planning. With proper planning, however, a system can certainly be developed which knows the location of every tagged item in the store, and can detect when items are mis-shelved or need re-stocking. Removing the line of sight requirement means that readers can be installed in shopping carts, for example, to monitor their contents and tell the shopper the total cost of his purchases, or suggest complementary items. For example, the clothing designer Prada uses RFID tags exclusively in its New York store. The dressing rooms can detect which clothes a customer brings in and call up information about them, suggest complementary items, and even show video of models wearing the same items [6].

Depending on the items sold in a store there will be different cost/benefit points for different retailers. Some items are very susceptible to "customer switching", a situation where a customer finds a desired item out of stock and purchases a competing item (one with a low "switching cost"). Having made the switch, they are often likely to keep buying the competing item. A "smart shelf" which knows when stock is getting low can prevent this situation by alerting the store manager to replenish the stock. Radio tracking systems can also record which items are removed from shelves and then compare this record with the items scanned at the checkout to determine if products are stolen. A pilot project in the UK records the time that razor blades are removed from and replaced on a shelf. If the blades are not paid for, the timestamp can be used to call up relevant security camera footage to help identify the thief. Currently such an elaborate system makes sense only for items whose aggregate theft constitutes a large monetary loss, but as innovation and economies of scale drive down prices for these technologies, they will become more viable for small-ticket items.

4. Conclusion

This paper has provided an overview of several location technologies, and a comparison of two popular identity technologies. It has discussed the relationship between the two: they are sometimes almost inseparable, and at others completely independent. An awareness of this separation is essential for a proper understanding of the role of sensing in automation systems.

References

[1] Localizers: Executive summary. Technical report, Aetherwire and Location, Inc. http://www.aetherwire.com/CDROM/General/AWL/execsum.html.

[2] The Auto-ID Center, 2003. http://www.autoidcenter.org.

[3] Paramvir Bahl and Venkata N. Padmanabhan. RADAR: An in-building RF-based user location and tracking system. In INFOCOM (2), pages 775–784, 2000.

[4] Andy Harter, Andy Hopper, Pete Steggles, Andy Ward, and Paul Webster. The anatomy of a context-aware application. Technical report, AT&T Laboratories Cambridge, 1999.

[5] Hitachi. μ-chip Product Specification, 2003. http://www.hitachi.co.jp/Prod/mu-chip/p0004.html.

[6] IDEO. Case study of Prada Epicenter, 2003. http://www.ideo.com/casestudies/prada.asp?x=1.

[7] Intersense Inc. MiniTrax Datasheet. http://www.intersense.com/products/prec/is900 MiniTrax.pdf.

[8] John Krumm, Lyndsay Williams, and Greg Smith. Smartmovex on a graph - an inexpensive active badge tracker. http://research.microsoft.com/~jckrumm/Publications/badgepublish.pdf, 2003.

[9] Micheal Di Paolo, William Furr, Nhi Nguyen Hearn, and Kem Kyoo Tae. Rfid: Replacement or supplement to bar codes? Master's thesis, University of Colorado at Boulder, 2000.

[10] Nissanka B. Priyantha, Anit Chakraborty, and Hari Balakrishnan. The cricket location-support system. In Mobile Computing and Networking, pages 32–43, 2000.

[11] Tony Seideman. The history of the barcode. In Carol J. Amato, editor, Inside Out: The Wonders of Modern Technology. Smithmark Publishing, 1993. Book not available, but article reprinted by Lasco Fittings, Inc (http://www.lascofittings.com/supportcenter BarcodeHistory.asp), in turn reprinted from American Heritage of Invention and Technology, a quarterly magazine issued by American Heritage.

[12] Symbol Technologies. Se 1200lr long range scan engine datasheet. Symbol Technologies web site. ftp://symstore.longisland.com/Symstore/pdf/SE1200LRFinal.pdf.

[13] The Financial Times. La Poste increases its range of services with the 'followed letter', March 14, 2003.

[14] Trakus, inc., 2003. http://www.trakus.com/svc hardware.html.

[15] Roy Want, Andy Hopper, Veronica Falcao, and Jonathan Gibbons. The active badge location system. In ACM Transactions on Information Systems, pages 91–102, AT&T Laboratories Cambridge, January 1992.

[16] Jay Werb and Colin Lanzl. Designing a positioning systems for finding things inside. IEEE Spectrum, September 1998.