Enhancing Identity with Location

A Study of the Behavior and Augmentation of RFID in an Automation Context

Hugo Mallinson
Institute for Manufacturing
Cambridge University
September 1, 2003

Acknowledgements

First of all I owe a great debt of thanks to Steve Hodges and Duncan McFarlane, who took me on at very short notice and gave me the chance to work on a project which I have greatly enjoyed. Steve's advice and guidance have been a great help, and his assistance during the preparation of this thesis was extremely generous. James Brusey had a few gentle suggestions which saved me untold frustration. Alan Thorne was very generous with his time as I muddled about with his robots. Simon Smith and Charles Crouchman solved innumerable little problems and always knew where to find whatever I needed. I would also like to thank my friends at the Cambridge University Boat Club, who have made this year a great pleasure. It has been a privilege to row with them. Katy has alternately cheered and prodded. I can't thank her enough. Most importantly I have to thank my parents, who have never doubted me.

Contents

Chapter 1

Introduction

1.1 Automated Environments

The capability to automate manufacturing has made a huge 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 thesis will refer to the differences between identity and location and 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.

1.3 Direction

This thesis focuses on a particular technology, known as Radio Frequency Identification (RFID), as a mechanism for providing both location and identity sensing in an automated system/environment. The thesis will demonstrate the potential and the limitations of RFID, and demonstrate how it could be augmented to provide better location information. It begins in the next chapter with a review of the various technologies available for determining objects' location. Chapter 3 will discuss technologies for determining identity, and will then explain the convergence of identity and location. This background will put RFID in the context of other location and identity technologies and allow the better demonstration of its advantages and its possibilities for augmentation. Chapter 4 contains a discussion of an empirical analysis of an RFID system, conducted to investigate the quality of its operation in a varied environment. Chapter 5 contains an analysis of these data and draws conclusions about what they indicate for RFID in various contexts. Chapter 6 discusses the possibilities for enhancing RFID. It discusses possible uses of the technologies from Chapter 2 informed by the results of Chapter 5. Chapter 7 indicates future directions for this research.

Chapter 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 chapter 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. These technologies will lay the groundwork for the analysis made in Chapter 5 and the recommendations in Chapter 6.

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 its 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

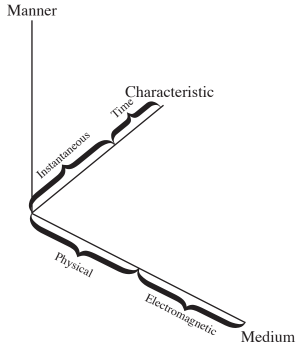

To give order and provide a basis for comparison of location technologies, a classification system was devised. The technologies can be categorized by three independent attributes: the physical medium through which they work, the characteristics of the medium which they measure, and the manner in which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

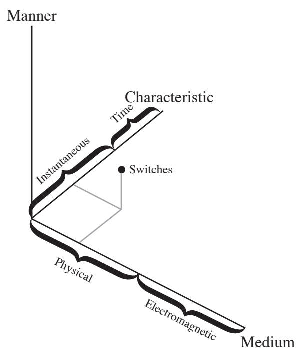

Switches fall very near the origin of the three-axis model: physical medium, low-complexity characteristics, low-complexity processing. From this point there are three ways to enhance this sensor. The complexity of the processing could be changed, for example by adding a second switch a small distance 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 well-placed 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 [33].

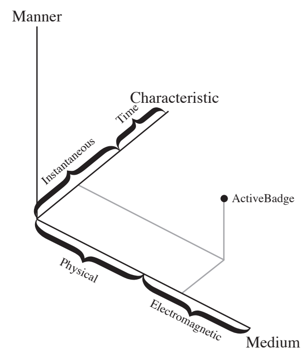

2.4 Signal Strength

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

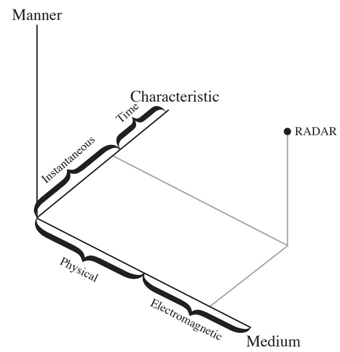

The RADAR system, developed by Microsoft Research, embraces this deviation of signal strength from the theoretical rather than trying to minimize it. The system consists of a number of 802.11b wireless ethernet base units arrayed around the space being instrumented. The item being tracked, in this case 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 50^{th} 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 [5].

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 [31]. 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: (a) synchronizing the timing of the transmitting and receiving nodes and finding the difference between the transmit and receive times or (b) 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-to-many 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.

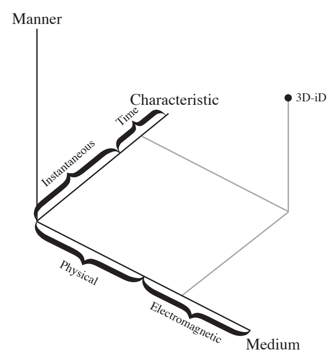

Ultra-Wideband

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 spread-spectrum 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 [3].

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 [3].

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

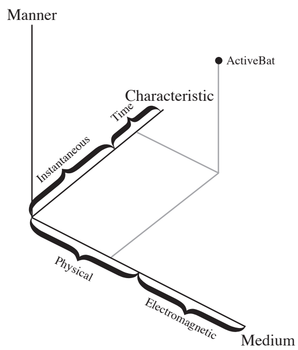

The ActiveBat was developed by Olivetti Research to provide greater accuracy than the ActiveBadge. The transmitter is smaller than the ActiveBadge, and the protocol is more complex. It consists of a wearable transceiver (5x3x2cm, 35g) and a network of nodes designed to be installed above a dropped ceiling. 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 [7].

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 [7].

The Cricket

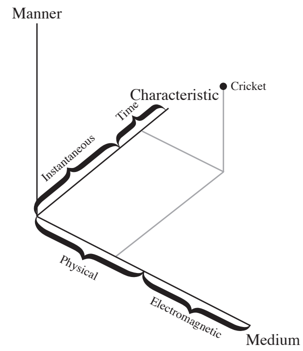

The Cricket system developed by the Networks and Mobile Systems group at MIT's Lab for Computer Science uses the same RF/ultrasound technique as the AT&T Bat, but with a decentralized topology. The Cricket system has no central processor which locates its nodes, rather, each node locates itself in terms 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 [20].

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 [37].

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 which have 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 [15].

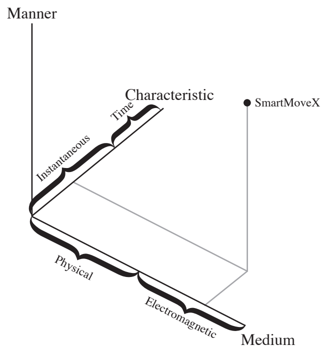

2.6.3 Inertial Tracking

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 [13].

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 [13].

Chapter 3

Identity Sensing

This chapter examines two leading technologies for providing identity data: barcodes and RFID. Combined with location data provided by the technologies in the previous chapter 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 [19].

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[30]), 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 [22]. 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 [26]. 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

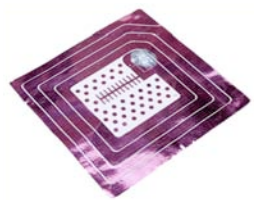

Figure 3.1: An RFID tag

Increasingly, automated environments are using RFID systems to provide location and identity information [19]. 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. An RFID tag (figure 3.1) consists of an antenna and a microchip mounted together in a substrate which is attached to the item being tracked. The RFID reader (figure 3.2) 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

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

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[9].

RFID systems are primarily used in one of four frequency ranges, as explained in table 3.1.

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

Figure 3.2: An RFID reader

[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.

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 $\frac{1}{d^3}$ – 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 countries and want to standardize their equipment, or for goods which are tagged (as in section 3.3) and then shipped internationally. These tags are also inductively powered and are limited by 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 3.1: RFID tag comparison

location data.

RFID technology is increasing popular in a variety of applications. The Auto-ID Center[4] 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.

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. In

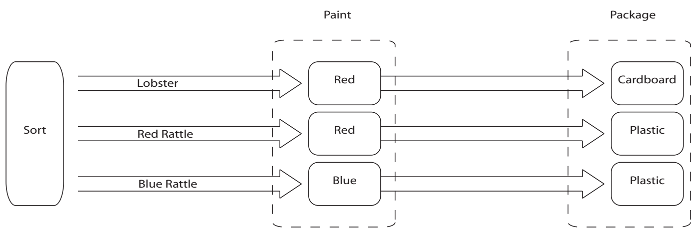

Figure 3.3: Factory A

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.

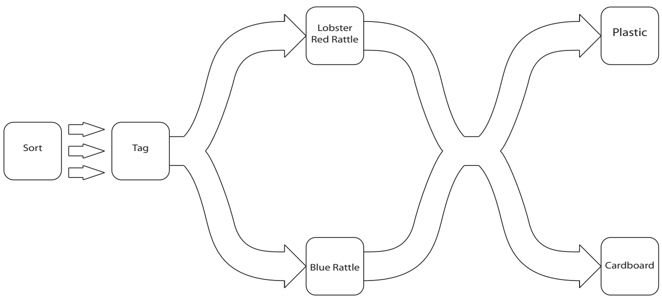

Figure 3.4: Factory B

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.

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 step 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 [12].

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.

Chapter 4

Investigating RFID

Chapters 2 and 3 provided a review of the technologies used in sensing location and identity. Chapter 3 then explained how these two technologies come together to allow objects to be tracked and gave examples of the usefulness of tracking. This chapter further explains the decision to examine RFID, and then discusses whether research could be conducted to provide still more information about the technology. Possible research is proposed to specifically investigate certain attributes of the tag and reader. This is followed by an explanation of the methodology used in collecting the data. Finally, the details of the data analysis are explained, in preparation for chapter 5 which will detail the findings.

4.1 Why look at RFID?

RFID is a mature technology, which has been in use for almost 20 years. Originally used for livestock tracking, it was then integrated into electronic toll-collection systems [16]. Further miniaturization of the tags made it possible to use RFID in access control applications. The technology is a very good fit for these applications, because it allows unique identification and therefore greater security than ordinary locks or keycodes – and is more robust than magnetic swipe cards – while not being significantly more expensive. But there is now a growing interest in using RFID for asset tracking [34]. This field has the potential to multiply the demand for RFID equipment many times over. There are retailers looking into attaching an RFID tag to every item they sell, using it as a radio version of the UPC tag. Major retailers Tescos and Wal-Mart have announced plans to use RFID to track goods in their stores. With this growing interest, RFID looks likely to become an extremely important technology in the near future [1].

One major barrier to wider adoption of RFID technology is the cost per tag.

As individual tag costs drop manufacturers can afford to tag more items; a lower price would make many more manufacturers consider the technology. Even at current prices – where it is often not economical to attach a tag to every item – there can be an advantage to tagging pallets or carriers in factories to monitor lots of manufactured goods [18]. While this may be cost effective for manufacturers, the real takeoff in tag usage will be when the price reaches a level where individual items can be tagged. If individual items are tagged by manufacturers the tags can be used at many points in the product's lifecycle, not just in manufacturing. A tagged pallet in a factory is only useful for the manufacturer but a tag installed in an individual item lets the item be tracked by the manufacturer, the shipper, the retailer, the consumer, and the waste disposal firm. These downstream users of the technology can join with manufacturers to demand innovation from the tag manufacturers to produce cheaper tags. The necessary price point for widespread adoption will eventually be reached with innovation by RFID manufacturers and growing economies of scale as more companies adopt the technology. Because the installed base is growing very rapidly, demand and pressure on these manufacturers to innovate is as well. This already-growing technology will take off exponentially when the natural course of innovation brings the price within the range of the largest customers.

4.2 Pros and Cons

Given the breadth of options presented in the last chapter RFID has some advantages and some disadvantages over other technologies. The tags are relatively cheap, rugged, and long-lasting, especially passive tags which don't require batteries. They avoid the major problems of barcodes: required line-of-sight and short range, and they can provide the flexibility of a writable tag. The price difference in the tags is what truly drives the divergence in price between the technologies as quantities rise, since barcodes can be implemented with only one-time costs (redesign of labels) and RFID tags can easily cost 20¢ apiece [14]. From an installation and troubleshooting standpoint RFID's operation is also not as intuitive as a barcode scanner, which shines an obvious beam of light on the tag being read. An RFID reading has to be taken on faith since there is no outward indication that a given tag is being read.

Overall though, RFID appears increasingly to be the technology on which many automated environments will be based, however it has not been studied empirically at much length in a manufacturing context. There is a great value to measuring more accurately how it behaves in a controlled environment so that the technology's limits can be better understood. This understanding will allow for more effective future deployment of the technology.

4.3 What this Research Will Show

The goal of this investigation of RFID is to show how the tag and reader interact and what location information can drawn from that knowledge. First the capabilities of RFID must be understood in an industrial context. This research studies the performance of RFID in the face of the most usual problems which appear in an automated manufacturing system. These are:

- A tag repeatedly arriving at a given point and being sensed (or not) by the reader. How likely is the reader to generate false positives and/or false negatives?
- A tag presenting a different profile to the reader, i.e. what happens when an item is not square to the reader?
- A tag of varying size. What difference does a larger or smaller tag make to the reader?
- Obstructions in the read zone. Physical obstructions between the tag and the reader, or simply near the environment.
- Transmission collisions with other tags. How does the system's response change with multiple tags in the reader zone. How about when two tags are very close to each other and could potentially affect each other's ID transmission?

4.3.1 The Possibility for Study

Can these situations be studied further? It would be possible to find cases where each of these problems occurs in industry and perform case studies of each environment. Then each scenario could be compared with similar ones in other environments and conclusions could be drawn. This approach would require finding appropriate sites, and conducting a series of interviews. It would have the advantage of providing data from actual manufacturing sites, but each site would likely be slightly different and this would have to be accounted for in the analysis. It would be very difficult to perform a side-by-side comparison of two different scenarios from two different environments, for example comparing the obstructed readers found in a shoe factory with the unobstructed ones found in a drink bottling plant. Variations could include the implementation of RFID, the layout of the environment, and the controlling software, just to name a few. Also, it would only be possible to study the specific issue present in each environment. If a question came up during the research and needed a certain configuration to answer it, another site which used this configuration would have to be found on very short notice.

Another way to study these situations would be to simulate them. A simulation is not the real thing, but it can be made as close as its designers wish. If an environment can be created where an RFID system can be tested in many configurations it would greatly simplify the research, as well as providing a test bed for future work. A simulation can more easily control for outside influences; since the people operating it have complete control and do not have to meet a production goal or run certain machinery at certain times they are free to shut down offending equipment or make many modifications to the environment not possible in a production setting. Most importantly, a simulated environment can have as many measurements made and controls installed as are required to capture the requisite data.

4.3.2 The Study Conducted for this Reserach

A simulation was created as a test bed for an RFID system. An RFID reader was installed within reach of a robotic arm, which was used to manipulate the RFID tag. The arm could place the tag anywhere about the reader, moving it in any direction and at any desired speed[1]. The lab which contained the robot also contained an electronic shuttle system and various pieces of equipment, but the reader was mounted as high as possible above the ground and far from other equipment so as to minimize interference from these surroundings. All surrounding electronic devices (mostly other RFID readers) were unplugged during testing.

4.4 Data Gathering

Simulating these scenarios is generally straightforward, but it is important to precisely state what is being measured. Tag collisions are particularly subtle since processing can happen in the tag, the reader, and even the serial line coming out of the reader and it is important to understand exactly which is happening so that the resultant data can be accurately analyzed.

4.5 Methodology

This study was conducted with RFID tags and readers manufactured by Checkpoint Systems. This equipment operates at 13.56MHz, the most common frequency for currently installed industrial systems [28]. While some newer installations are moving to the UHF frequency bands, there are still many new systems

[1] The robotic arm moves at a maximum of 2 m/s, and for safety reasons is rarely run at more than 30% speed.

being installed at 13.56MHz, in addition to the many existing installations [27]. The data were collected by moving the RFID tag in a set lattice of points around

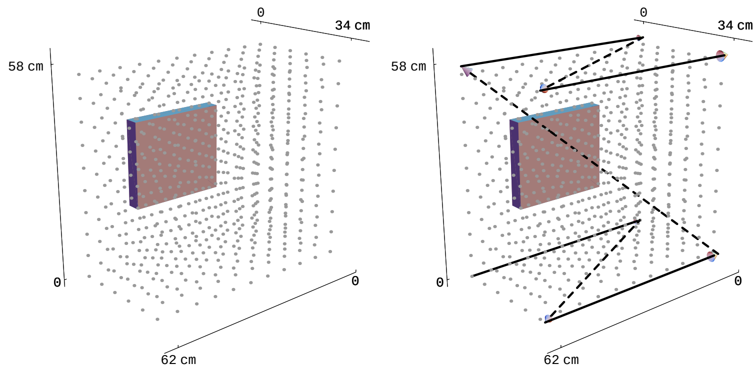

(a) The lattice of points where measurements are taken

(b) The path which the tag, attached to the robot, follows. Data are taken while it moves along the solid lines, the dashed lines are movement without data gathering.

Figure 4.1: The data gathering region

a Checkpoint Slimline RFID reader. The reader was placed so that any physical obstructions in the environment were out of its range. The metalwork which fixed it in place was attached to the back of the reader, and in the same manner as would be used to install it in an industrial environment. The tag's position was incremented first horizontally across the reader, then away from it, then vertically up, as shown in figure 4.1. The tag was held in position for a set time interval, and the reader was monitored to see how many times it successfully read the tag in that interval. The reader was first monitored to ensure a consistent read rate, so that the number of reads could be said to be proportionate to the read rate. While the output tended to come in bursts, over a period of a few seconds they average out and the 5 second dwell time used in the data gathering provided read numbers repeatable to within 1 or 2 reads out of 45 in initial testing (repeatability will be fully discussed in the following chapter), or an error of a little over 2%.

The tag is manipulated about the reader using a Fanuc M6i industrial robot arm. The use of a robot guarantees that positions can be repeated to sub-millimeter accuracy over as many repetitions as are necessary for gathering the data. The robot can return to any position with an error of ±.08mm [21]. The data are collected on 2cm and 4cm centers, so this error constitutes at

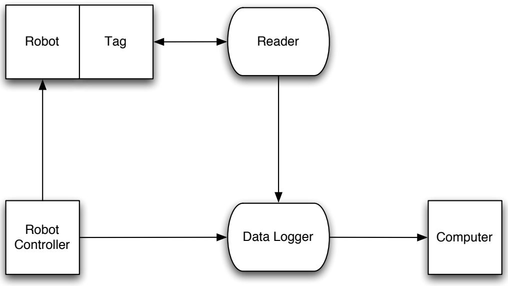

Figure 4.2: The data gathering process

most 0.4% of the distance and can be disregarded. The robot controller is also capable of inserting the time delays necessary to properly collect the data. An I/O line from the robot connects to a custom data logger, which is activated for a 5 second burst once the robot has moved the tag into its next position. The data logger is a PIC microcontroller activated by the robot controller when the tag has been moved into position. It receives the RFID reader's serial output and transmits to a PC. When it is active it counts the number of times the reader successfully interrogates the tag. On deactivation it sends this total number of reads along with an index number to the logging PC. The PC collects these logged data to a file, which is then processed to return the spatial dimensions to the serialized data and store them in matrix form. It is then an accurate depiction of the testing environment and can be manipulated without losing any of the spatial relationships between the points. The robot control files are generated on the computer as well, which gives the maximum flexibility in their implementation. They can be generated to record data points at any interval along any of the three axes and in any orientation. Measuring the read quality of the reader and tag could be done in a variety of ways. Potentially the strength of the EM field itself could be measured, then correlations between field strength and tag reads could be made. Since this thesis focuses on the implementation of a specific tag technology and manufacturer, the goal was to strike a balance between testing in a complete vacuum and working with the actual industrial equipment as it would be used in a factory. This balance was best met by choosing to measure the number of times the reader successfully read a tag in a set time interval as the metric for evaluating the "quality" of a read.

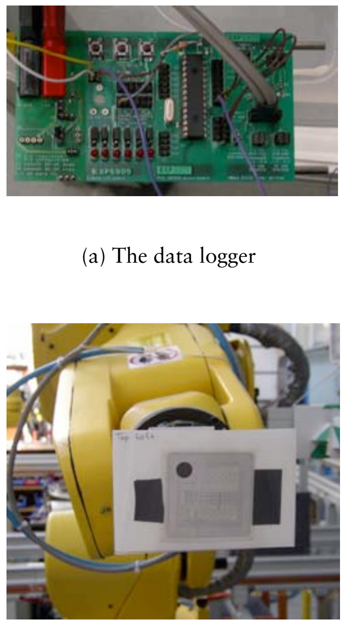

(a) The data logger

(b) The 2" Checkpoint tag attached to the robot

Figure 4.3: The experimental setup

The Checkpoint tags used in this research use the MCRF355 chip from Microchip technologies [24]. When activated, the MCRF355

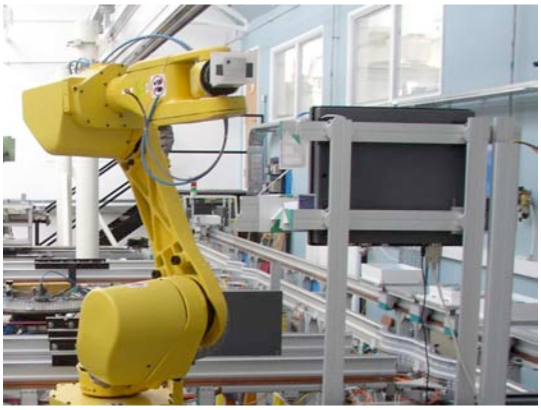

Figure 4.4: The Fanuc M6i and Checkpoint reader in the lab environment

transmits its 154 bit string on the 13.56MHz band, modulated at 70kHz, therefore taking 2.2ms per transmission [17]. It then pauses for 100ms before transmitting again. Thus in a 5 second interval the maximum number of reads from a single tag is roughly 49. The 5 second interval was chosen after it was determined to be long enough to allow for variances in the reader's output (serial port buffering, etc.) to normalize, and short enough to make the collection of over 15,000 data points feasible without undue changes to the environment.

Chapter 5

The RFID Reader Data

This chapter documents the data collected to further investigate the Checkpoint readers and tags. It explains and demonstrates the various methods available to best visualize the data, and draws conclusions about the various sets collected.

5.1 The Presentation of the Data

There are many possible ways to present the data gathered in this thesis. Since each dataset contains thousands of points, it was very important to find ways to make the sets meaningful both visually and numerically. Figure 5.1 is an excellent example: it contains two different representations of the number of reads data, each with a specific value. The first is a density field rendering of

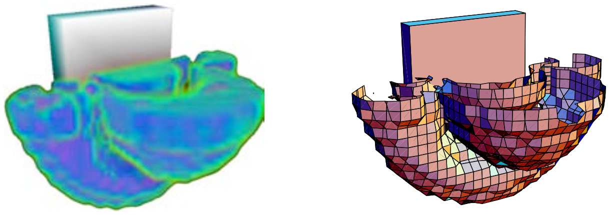

Figure 5.1: The RFID reader field

the area around the reader, cut away to show the cross section. The second is a contoured density plot, showing a contour, in this case the boundary with 0 which shows the outermost shape of the read field as well as the dead spot in the middle. These two graphs are three-dimensional renderings of the cross sections shown in figure 5.2. These plots give the most precise rendering of each point, while still giving a sense for the overall shape. Another interesting rendering is the difference between datasets. Figure 5.4 shows the distribution of differences

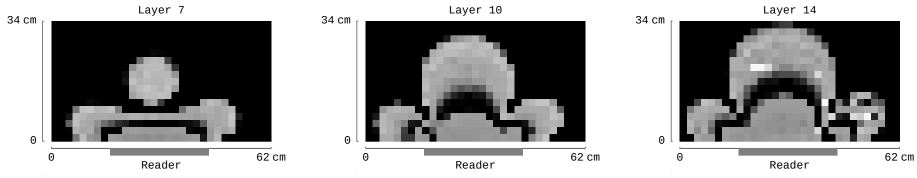

Figure 5.2: Three cross sections from the data renderings in figure 5.1

between the first and second sets of data taken with the initial configuration.

5.2 Assessing the Experimental Setup

A large portion of this research was involved with establishing a proper methodology for collecting data. Once data could be easily logged it was important to establish that the experimental setup gave useful, repeatable results.

5.2.1 The Initial Dataset: Establishing a Baseline

The first dataset (Appendix A.1) was collected at 2cm intervals in a lattice of points at whose outer limits the tag was out of the reader's range. These data reveal a regular shape to the reader field, shown in figure 5.1 in section with the reader in the background. Of immediate note are two important attributes. First, the pattern of the read field is not an ellipsoid, but rather lobed, with a sharply varying outer limit. The reader's range is significantly shorter in the area of the notches between lobes than their surroundings would cause the casual user to expect. Secondly there is a dead area about half way from the reader to the outer extent of the read range.

These 'unexpected' dead areas could prove very useful for processing algorithms to enhance the location resolution of the reader. They provide an isolated area much smaller than the range of the full reader which could be harnessed to locate tags, for example if the reader were mobile and moved only slightly it could move a tag in and out of the dead area, and thus enhance the location resolution.

5.2.2 Confirming the Initial Data

Once these data were collected, a second set of data was collected with the same apparatus to examine any fluctuations in the experiment which were not controlled for. These data were compared with the initial set, and are shown in figure 5.3. The shape of the data was similar, still symmetric and free from random fluctuations. The distribution of the differences between the sets is shown in figure 5.4. A difference of 0 indicates a perfect match between the sets, and these data are tightly grouped about 0, with 87% falling within one standard deviation. The consistency of the pattern of reads and the tight grouping of the difference data suggest that this pattern is highly repeatable.

(a) (b)

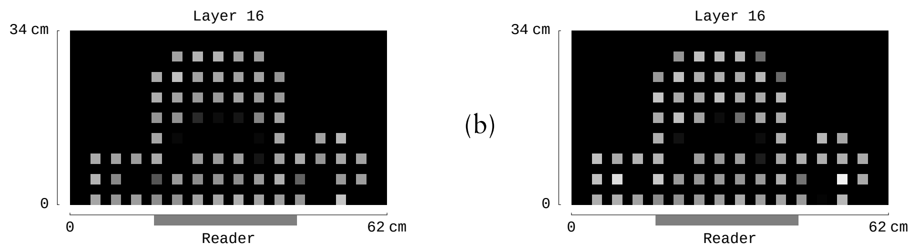

(c)

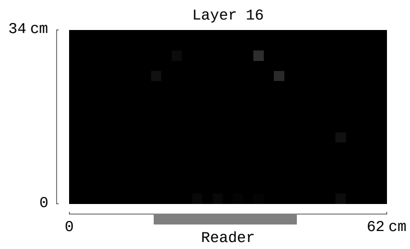

(a) The original data, only points taken in the second set are plotted
(b) The second set with an identical configuration
(c) The difference between the two sets

Figure 5.3: Confirming the experiment's repeatability: two datasets taken with an identical configuration

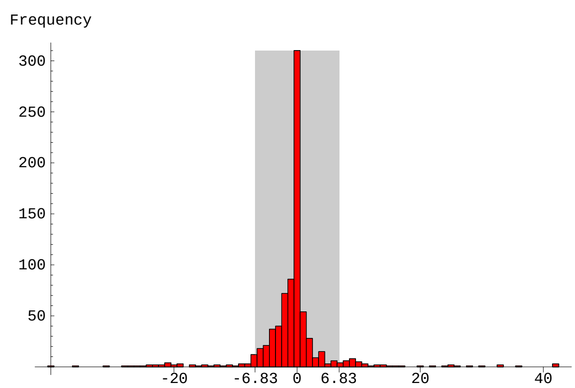

Figure 5.4: The distribution of the differences in number of reads between the first and second datasets. 87% of the data, highlighted in gray, are within 1 standard deviation, 6.38, of 0

Point	Standard Deviation	Mean	% within x standard deviations of the mean		x% of the data are within mean $\pm$	
			1	2	90%	95%
9	1.07984	39.84	54%	99%	2	2
7	1.22268	38.80	76%	97%	3	3
8	1.30775	38.63	55%	94%	3	3
6	2.13948	38.22	72%	95%	4	5
2	2.77099	20.72	73%	98%	5	6
1	3.21579	25.39	68%	98%	6	6
4	4.83982	32.52	88%	95%	6	7
5	4.69484	30.67	89%	97%	6	9
3	4.28381	36.75	89%	93%	5	14

Table 5.1: The repeatability statistics for a sample of 9 points

5.2.3 Confirming Repeatability: Multiple Reads

To further measure the repeatability of the apparatus, a sample of 9 points were measured 100 times each in quick succession. The points which had the very highest mean number of reads also had small standard deviations, indicating a high repeatability (figure 5.5(c)). Those with lower means tended to have a larger standard deviation, but it can be seen in figures 5.5(a) and 5.5(b) that there is great variation, with standard deviations as high as 4.8. 4.8 is still quite small compared with the possible number of reads (roughly 50). These numbers are summarized in table 5.2.3. If the readings are distributed with a Normal distribution, then 68% of the data are within one and 95% are within 2 standard deviations of the mean. Almost all of the points are actually more tightly clustered than this. In general

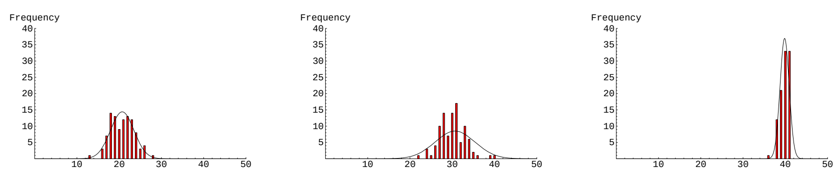

(a) (b) (c)

Figure 5.5: 100 samples of the number of reads at 3 different points

the points with higher numbers of reads tend to produce lopsided plots, as in figure 5.5(c). Points with lower average rates show a more symmetrical distribution tailing off at either end, suggesting that the higher points which don't tail off do so because they hit some sort of threshold or maximum read rate. These results could also be explained by the power provided to the tags. The graphs with more variation (larger standard deviations), such as figure 5.5(a), generally had lower numbers of reads also, well below the theoretical maximum of 49 calculated in 4.5. Since the tag did not achieve this number of reads, something must have caused it to not transmit during some of the 5 second interval. There are two parts of the transmission process which could fail: the powering of the tag by the reader, or the transmission of the tag data back to the reader. It is possible that the tag lost power during some of this period, as would happen if the power level were fluctuating near the minimum threshold the tag needs to run. Such a fluctuation would explain both the lower average, since the tag would not be fully powered during the 5 second cycle, and also the wider standard deviation in the data, since the difference between the theoretical upper bound and the mean is greater in these dataset than in set 5.5(c). It is also possible that the tag was powered, but its output was insufficient to reach above the noise floor and be detected by the reader. This could be either a function of the tag or of some source of noise in the environment. Since a range of values are found in all plots, and they appear in a regular pattern it is unlikely that the sub-maximal transmission is due to an external noise source.

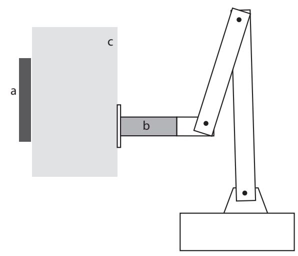

(a) The long end effector

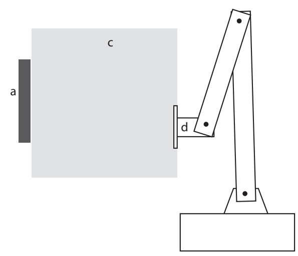

(b) The short end effector

Figure 5.6: The robot at its furthest from the reader. The long end effector (b) keeps the tag away from the robot, but also forces it much closer to the reader (a), cutting significantly into the read zone (c) provided by the short end effector (d)

5.2.4 Checking for Systematic Errors

Restrictions in the experimental space meant that the robot could not be positioned ideally with respect to the reader, so some compromises had to be made to allow the data to be gathered. The end effector used to attach the tag to the robot arm was 12cm to reduce unnecessary flex and assure the precision of the tag's placement, and also to allow the robot to position the tag in all the necessary locations. The proximity of the robot – a large ferromagnetic object – to the tag could have an impact on the readings. Since the tag both receives its power from and broadcasts by modulating a magnetic field this could be a factor. A significantly longer end effector was constructed to examine the differences between end effectors. At a length of 38cm it was nearly four times the length of the 10cm end effector.

A dataset was taken with the long end effectors and compared with the original dataset taken with the short end effector. The

experiment requires that the robot not come between the tag and
the reader, but the robot's configuration does not allow it to move far enough away from the reader to reach the farthest points reached with the short end effector, as shown in figure 5.6(a). As a result only 800 data points could be taken where the short effector was able to reach 1890 – completely covering the range of the reader. Examined visually, the results from the long end effector are less regular. They also do not show the symmetry of the two previous short effector datasets. Two changes were made in the configuration in using the longer end effector which might have caused the disruption. Because of the different size of the longer effector, the configuration of the robot's joints were different for each point. To put the tag in position, the robot end of the long effector was much closer to the robot than with the small effector, this meant the robot's second joint was generally closer to the tag than with the short effector. Each joint contains a large DC motor and associate circuitry which could have introduced new disruptions to the magnetic field. Secondly, the longer end effector had to be reinforced with large aluminum grommets to prevent wear. Although only 1.5cm across, the introduction of new metal to the environment could affect the magnetic fields.

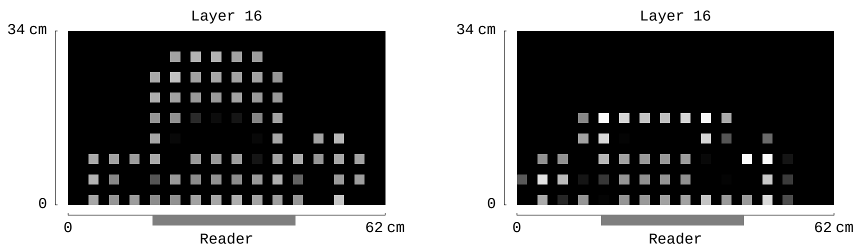

(a) Data with short end effector (b) Data with long end effector

Figure 5.7: The effect of the long end effector

The long end effector data did not resemble as closely the baseline data as did that of the short end effector, and while the longer effector does move the tag farther from the robot and produces different data, it also severely limits the range of measurement. The possibility that using the longer end effector actually introduced some kind of interference into the environment cannot be disregarded. With the high correlation between the first two short effector sets, subsequent data can be judged against them as well and conclusions drawn since any effects from the short effector will be regular across the data.

5.3 Testing Potential Disturbances

5.3.1 Angled Data

In order to test the response of the reader to irregular tag placements, data was collected with the tag rotated 45° through its vertical axis. For comparison with the other data sets, the positions recorded were those of the center of the tag. The results can be best shown by comparing two horizontal slices of data. The first, figure 5.8(a), is the baseline data. The second, figure 5.8(b), is the data collected with the angled tag. (Note that the line of data closest to the reader could not be collected with the angled tag). While the initial data has three clear lobes, the angled data is almost completely missing one lobe. Additionally, the points in the side lobe still present tend to indicate higher read rates. although there are few new points, that is points which had read 0 before, rather some points with lower read rates in the parallel data show higher rates in the angled data. These results are consistent with the nature of the tags. The Checkpoint

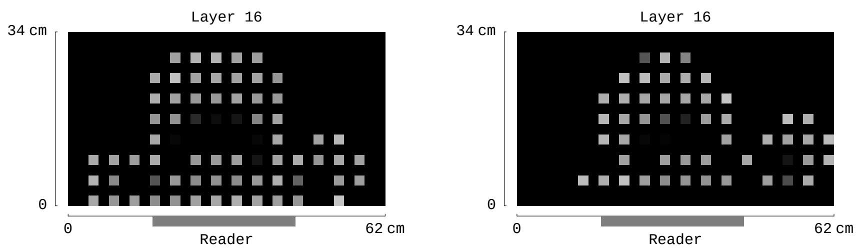

(a) Tag parallel to the reader

(b) Tag at 45° to the reader

Figure 5.8: The results of changing the tag's attitude to the reader (the row of points closest to the reader cannot be read with the angled tag)

tags are inductively powered, that is they draw the power for their circuitry from a magnetic field generated by the reader. If the reader is modelled as a coil of current-carrying wire (an approximation of the actual internal antenna) the magnetic field lines propagation through it can be see to fan out across the read zone, shown in figure 5.9. The tag receives the greatest power when these field lines are perpendicular to its antenna and it presents its greatest cross-sectional area, and least when they are parallel and its profile is negligible. As can be seen clearly in this diagram, the angled tag presents a larger profile to the magnetic field on one side of the reader than the other, and this could very well account for the variation in read rates. With the tag angled the reads stop abruptly at one edge of the reader (the "unfavoured" side) and continue beyond the reader on the other. The readers are designed to read tags at $\pm 40°$, so there is no problem reading points in the middle of the field, as seen here [25].

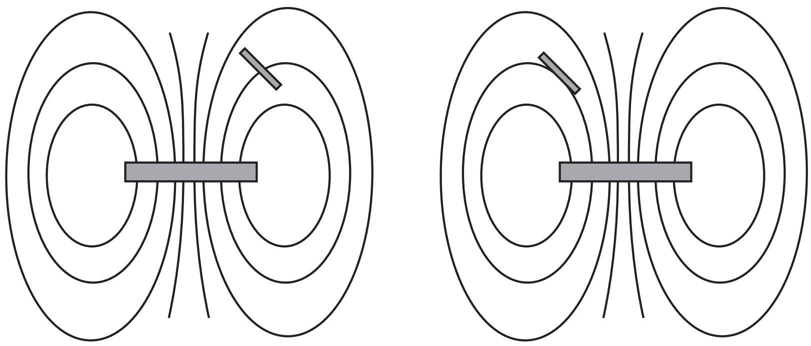

(a) The tag face perpendicular to the reader's magnetic field. This position will induce the maximum current in the tag's antenna

(b) The tag face parallel to the reader's magnetic field. This position will induce the minimum current in the tag's antenna

Figure 5.9: The reader with its generated magnetic field lines and the tag

5.3.2 Larger Tags

A set of data was taken using a larger Checkpoint tag [24]. This tag uses the same Microchip IC, but in a format with 4 times the antenna surface area. The larger tag produced the same field pattern as the original tags, but there was less variation in the number of reads. The graphs in figure 5.10 show only a few different read rates with a much sharper drop-off at the edges of the read zone. The dead spots in the middle of the reader are still present, although they are smaller.

That the larger tags provide consistently high numbers of reads with less tailing off to low numbers is not surprising, since their larger cross-sectional area means more magnetic field flux through the tag and so more power delivered to the microchip.

The larger tag provides a more sharply-defined field with its tighter clustering of read quantities. The graph in figure 5.11 shows a sharp drop above 54, while the small tag tails off more gradually and gets reads at a consistent level from roughly 30 right down to 1. For applications where the larger size of this tag can be accommodated, the sharper read field should be a boon for system designers as the sharp transition from 0 reads to maximum found at the outer boundaries is easier to detect than the more gradual form of the small tag. The distribution

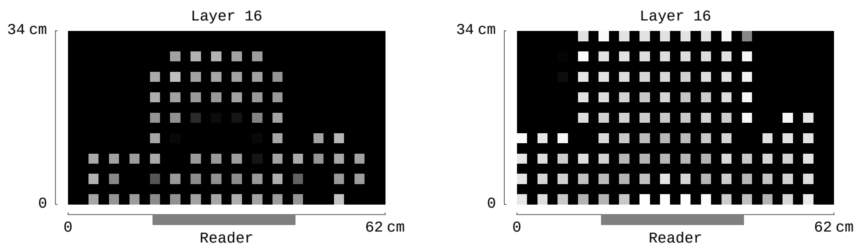

(a) The original data

(b) The same section using a tag with 4 times the cross-sectional area

Figure 5.10: The response of the reader to differently sized tags

of the number of reads moves significantly for the larger tag as well, as seen in figure 5.11.

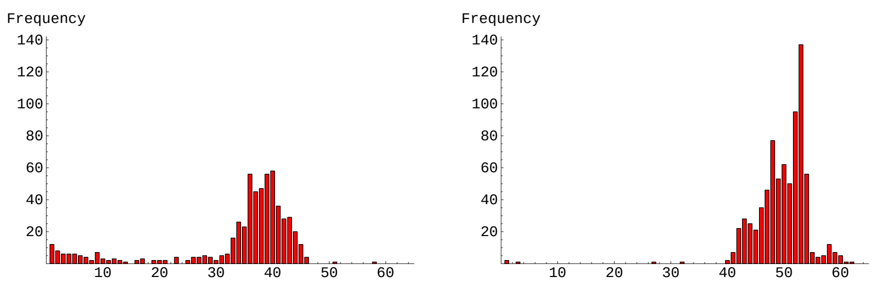

(a) The small tag

(b) The large tag

Figure 5.11: The number of reads with large and small tags. The 0 frequency is not included to highlight the shape of the read points.

5.3.3 Tag Collision

In many industrial environments there will be more than one tag which appear within range of the reader at the same time. All these tags will be powered and will attempt to transmit. If two tags transmit at the same time the reader will not be able to read them – it is not possible to differentiate one signal from the other without using a coding method such as CDMA[1] or ALOHA [23].

Tag collision is an extremely important factor in RFID system design: in a given environment, as the reader and tag range increase more tags will be in range and the likelihood that two tags will share a reader will increase as well. The Checkpoint tags use a simple anti-collision system to avoid transmitting over each other; it could be described as "passive." Each tag transmits and then sleeps for 100ms before transmitting again. The timers used in the tag microchips are intentionally imperfect, so the delay varies from tag to tag and over time the transmission times will drift with respect to each other. These random transmission times mean that many tags can be read while sharing the reader's space. Checkpoint calculates that 25 tags can be read in one second [24].

Figure 5.12: The second tag installed next to the reader, providing a constant stream of transmissions to test the tags' anti-collision performance.

To determine what difference multiple tags made to the reader a second tag was placed within the reader's range while a dataset was collected. In the first experiment the tag was placed directly adjacent to the reader, where it would be in range throughout the data gathering and in the very highest read rate zone. Since the tags broadcast for a very short time period (2.2ms) over their duty cycle (around 100ms), it would be logical to expect there to be little difference in the results, since the probability of two tags transmitting at the same time should be around $\frac{2 \cdot .0022}{.1} = .044$. So for a given point the readings could be expected to be around 96% of their value without tag interference. These calculations assume a 100ms interval between transmissions and do not take into account any time delays for the reader to process the tag data. The data logger software was rewritten to differentiate between the two tags and record reads for the original tag as well as total reads, from which the reads for the new tag could be calculated. The shape of the data is similar to previous sets, as seen in figure 5.14. In the second experiment the second tag was placed adjacent to the tag, attached to the robot end effector so that it would sit

Figure 5.13: The second tag installed adjacent to the first tag.

[1]Code Division Multiple Access is a method for dividing bandwidth among many users. Transmissions are encoded according to a code assigned by a central server, and the encoding allows many transmissions to take place simultaneously while still being intelligible to the receiver. CDMA has achieved great popularity as the predominant protocol for mobile telephones in North America.

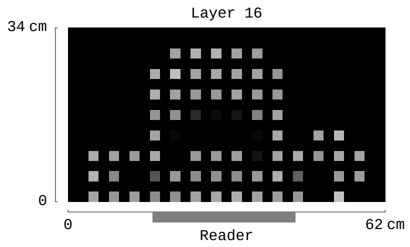

(a) The original data

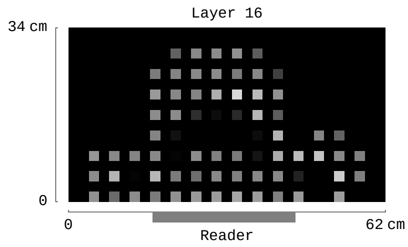

(b) The same section with another tag directly adjacent to the first

Figure 5.14: The response of the reader to collision between non-adjacent tags

in a similar position to the first tag with respect to the reader for all readings. In the first experiment the pattern is unchanged, but the number of reads falls. In the second, there is a slight change in the configuration of the points. The figures have become slightly asymmetrical. It is unclear what has caused this, since the experiment was performed under the same conditions as previous ones. There is potentially an interplay between the tags caused by their different heights: the tag whose data is being recorded is at the position indicated in the plots, but the other tag is mounted directly above (see figure 5.13) and so interacts with a slightly different reader field. Since the "interfering" tag sees a different field at every data point its transmissions will not provide the constant error seen in the previous experiment. It is possible that this difference in height, which effectively shifts the reader field 3cm down for the second tag, magnifies slight variations in field symmetry at different heights. Figure 5.16 shows the distribution of numbers of reads compared with a single tag, which shows the drop in average read rates apparent from the cross section plots.

5.3.4 Obstruction

Another important issue for RFID readers is the presence of obstructions. It is rare that a tag will be read in an empty room, generally there are conveyor belts, other products, people, shelves, and any manner of other potential obstructions. The Checkpoint tags operate at 13.56MHz, which is generally less susceptible to physical interference than the higher frequencies like 2.4GHz, especially from people. To investigate the changes in the RF field due to an obstruction, an aluminum bar 8cmx1.5cm was placed flush along the bottom of the reader face and a dataset was taken. One cross section is shown in figure 5.17(b). The bar was

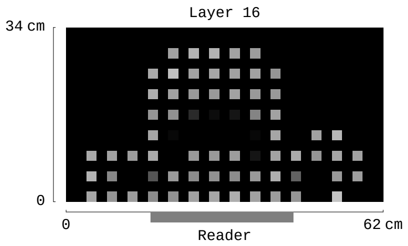

(a) The original data

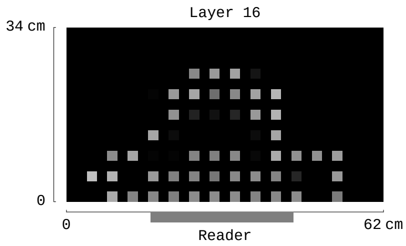

(b) The same section with another tag in the read range

Figure 5.15: The response of the reader to collision between adjacent tags

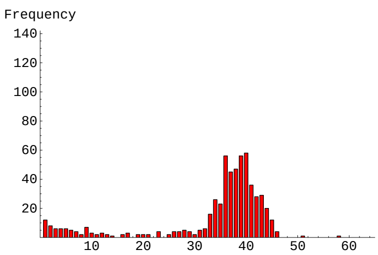

(a) Single tag

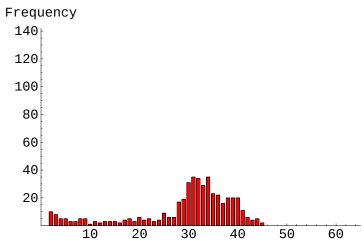

(b) Single tag with adjacent tag providing interference

Figure 5.16: The distribution of read numbers with and without interference

then moved 10cm out from the reader face and another dataset collected, shown in figure 5.17(c). As expected, the obstructions severely limit the interaction between tag and read. In the second case there do not appear to be ancillary effects to the pattern of reads, the data simply appear to have had a chunk removed around the area of the bar. In the first case, however, where the bar was placed directly in front of the reader, it appears to have greatly reduced the size of the read field, while still maintaining the lobed shape apparent in the other datasets. Figure 5.18 shows the position of the bars and gives a sense for the shape of the overall dataset for each of the three scenarios. It should be noted that no data was collected directly above and below the obstruction, so the lack of points in this region does not necessarily indicate that reads would not occur there. These renderings show very clearly the different appearance appearance of the read field caused by the difference in obstruction.

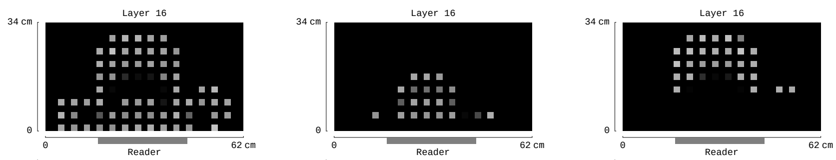

(a) The original data (b) The bar flush with the reader (c) The bar 10cm from the reader

Figure 5.17: The response of the reader to collision between adjacent tags

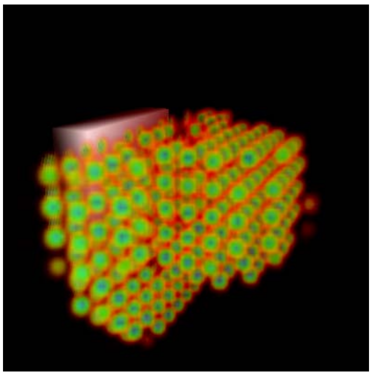

(a) The original data

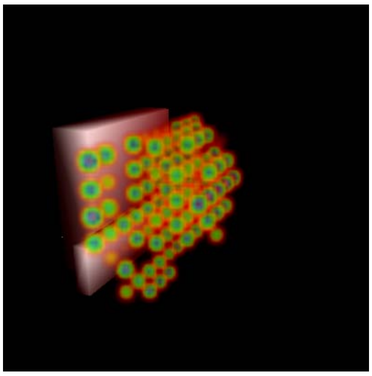

(b) The bar flush with the reader

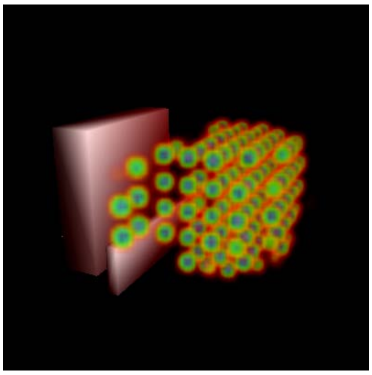

(c) The bar 10cm from the reader

Figure 5.18: The effect of obstructions on the shape of the read field

Chapter 6

Summary

6.1 What Has Been Shown

This thesis has investigated identity and location and their interdependence. It has looked at various technologies which provide both pieces of information. It has also examined some instances where these two data are used in concert. Finally it has investigated one specific technology, RFID, because of its growing importance and potential for further study, to better understand its potential for providing location data as well as its limitations. To this end a large amount of data was collected, testing various attributes of the technology, with the goal of establishing a picture of the range and suitability of RFID in various situations. The analysis of these data have produced two broad results:

- A clearer understanding of the range and repeatability of RFID, which permits a greater understanding of the behavior of RFID systems in an automated context.

- An understanding how RFID can be used for location information. These data show how RFID could be enhanced to provide what improvements could be made to RFID in terms of location. The analysis performed in this thesis indicates further potential for research to enhance the location data extracted from an RFID system. It also suggests situations where RFID would need to be augmented to achieve the level of performance appropriate to a given situation.

6.2 Application of the Research

The data sets discussed in the previous chapter provide a clearer understanding of where an RFID reader and tag best interact. The data show an area about 20cm in front of the reader where tags are read reliably at a high rate and without dead spots. In an environment with fixed readers; for example, reading objects which pass on a conveyor belt, it is important that the objects have the best chance possible of being read. From these data the optimum position can be found which will let the tag pass through the largest possible volume where reads are possible.

Another important fact learned from these data is that the read rate is not simply a function of distance from the reader. Faced with a non-working reader, intuition would tell most people to try moving the tag closer, but these data show that the area closer to the reader not only doesn't have a higher read rate, but it also has more dead spots and thus a lower overall likelihood of getting the tag read. The lobed nature of the read field is important to consider as well. It is more complex than a simple ellipsoid field and quite difficult to visualize accurately when examining an installed reader. The charts developed in this thesis could be profitably used during the installation and configuration of readers to guarantee that a point which appears to have good reads is not just a transient reading and does in fact lie in a high-reliability zone.

6.3 Enhancing RFID

6.3.1 Building on the Data

As seen in the last chapter, there is room for RFID itself to be enhanced to provide finer-grained or more reliable location data. The question is in which situations would enhancement be necessary, and if neccesary how to go about it. Producing finer-grained location data could allow a single reader to provide information for which a series of readers is now required (figure 6.1). Finer grained location would be most useful in areas where tags need to be tracked to a high degree of accuracy over a large area. Currently, readers are treated as binary sensors, so the resolution of the tag location information is proportional to the number of readers. If a reader were constructed to cover a larger area, the location information which it provided would be more detailed than that of the original reader. The most reliable tracking happens where the tags are confined to one dimension of movement, as on a conveyor belt. As soon as tagged items are loaded onto pallets and moved around a warehouse, for example, they become much harder to track. To sense location in two dimensions an array of readers with overlapping fields, as in figure 6.1, are required. Given data from

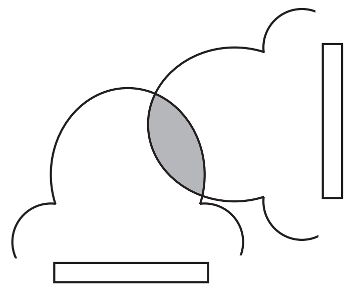

Figure 6.1: Two readers can be used in concert to enhance location information. If a tag is read by both readers it is in the gray region.

such a configuration, the calculations to locate a tag are non-trivial, as they require allowance for field overlap and possible dead spots.

Another possibility could be to have a mobile reader which sweeps across a given area. A mobile reader would then be directly above the tag it was reading, minimizing the distance to the tag and also interference with other readers. Such a system would require additional hardware to move the reader; a reader which can identify many separate locations would have a lower cost and fewer moving parts than a reader which has to be moved around. A single reader with precise sensing become increasingly valuable as the complexity of both the environment and with it the location sensing requirements increase.

6.3.2 Possible Enhancements

Processing

The number of reads density plots show a rapidly changing field around the reader which could be exploited to locate tags more precisely in the read field. A simple algorithm could move the reader through a preset range and then, based on the number of tag reads it found in each position, find the point in the read field where the tag was most likely to be (along the lines of the RADAR and SmartMoveX systems). While this seems cumbersome, there is no reason to believe that with some refinement a system could not be developed to move readers over only a small area in a well-chosen path. This would place tags more accurately than simply “in range” of the reader. If a reader could modify the behavior of its antenna internally, then it could automate this procedure: reading for a set time interval with one configuration and then changing to another and comparing the pattern of reads in each against a pre-determined lookup table to find the most likely tag positions.

Augmentation

It is also possible to consider augmenting the design of the RFID equipment itself. One possible modification would be the addition of Received Signal Strength Indicator (RSSI) circuitry which would return data about the strength of the signal received from the tag. These data could be used to estimate the distance from the tag to the reader, enhancing the position information. Alternatively, the amount of power sent to the tag could be varied until the threshold where the tag stopped being read was found. These data could be used in a similar manner.

With further modifications to the reader it could also be possible to perform Time of Flight calculations by sending pulses at known intervals to determine the distance to the tag. Most HF readers (like the Checkpoint) work in an "always on" format, detecting changes in received signal to generate reads, so this would require sending timed pulses to obtain proper ToF data. UHF spectrum requirements mean that UHF readers do not operate in an "always on" format, and so are already naturally suited to such a modification.

Additional Sensors

The last, and possibly most obvious way to enhance an RFID reading is to use additional sensors. In some situations it may simply not be feasible to use an RFID-only approach and the system may need the precision of break-beam or tactile sensors or a long range, as that provided by cellular telephone networks or the Global Positioning System[1] to guarantee object locations. In such a situation RFID can provide location data which is augmented by additional sensors, or it can be used solely for identity.

6.4 When RFID is not Enough

What these data have shown very clearly is that RFID is not perfect for all applications. It senses identity very well, and location quite well, especially if precise resolution is not required. The data bear out what instinct and experience with radio systems tell us: radio devices' performance always varies, and sometimes because of apparently minor disruptions. RFID will never be 100% repeatable, as the repeatability plots show. It is not an appropriate technology for saying absolutely where an object is. When location has to be known to the highest precision, it is practically impossible to improve on the accuracy provided by a physical stop which keeps an object precisely in place. For example, an assembly line for mobile phones a robot solders a microchip into certain phones. RFID could identify the passing phones and when it found one which needed the chip it could trigger a physical barrier which would hold the phone exactly in place, guaranteeing that the chip was installed in the proper position. With time and further research RFID's location resolution can likely be improved, and steps can be taken to enhance its repeatability, but there will always be situations where it will not offer adequate precision or accuracy for location sensing. This will not diminish its usefulness in industry, rather it will simply be applied to the situations in which it operates best.

[1] GPS is a technology which can locate a device to within 3m almost anywhere on earth. It uses high-precision timing and requires a clear view of at least 4 satellites so it is not appropriate for indoor use and was not discussed in this thesis.

Chapter 7

Conclusion

As a popular and growing hyrbid of location and identity technology rfid was chosen for an empirical examination. This thesis has explained the existing technologies in each category, and weighed the pros and cons of their use in an automated environment. This examination revealed the complexity of the system, but also the underlying patterns and great possibilities for the technology and its augmentation.

7.1 Future Work

This thesis has continued the ongoing discussion of work which could be done to better use RFID technology in industry. It has approached the problem from one angle: empirical testing of a specific setup in a controlled environment. The research and analysis indicated areas which should prove fruitful for future work;

- This research focused on the interaction between tag and reader. There was no work done on the inner workings of the RFID readers, and this area especially appears to hold promise for future innovation.
- Distance measures based on RSSI data could well prove useful and should be investigated more fully. There are also fundamental differences in how energy is sent to the tag and how the signal is received back from it, and these differences could bear further investigation.
- There are anecdotal problems with readers reading tags that they cannot normally because the tag is being powered by a second reader. These potential effects were eliminated in this research and therefore not studied.
- The protocols used by readers to communicate with the tag vary, and could

therefore have an effect on read rate data, as could limitations of the circuitry of both the tag and reader.

- The analysis in this project is based on read rates. Early research determined that reads happen, on average, at constant intervals over the measured time. While these project results bore out the conclusion that read rates with 13.56MHz readers are generally constant, it is possible that subtleties of circuitry and protocol also have an effect on the read rates and analysis still needs to be done with readers and reader antennas of other frequencies to determine an overall metric for describing RFID systems. The most striking differences will likely be between HF readers like the Checkpoint series, and UHF readers in the 2.4GHz range, because operating at a different frequency may produce very different patterns of read fields. In addition, the two different frequency ranges have different licensing requirements which result in different technologies being used in the tags.

- The different anti-collision techniques used at different frequencies will also play a large part in the response seen in multi-tag environments. Tags like the Checkpoint with a relatively 'dumb' anti-collision method ought to respond differently from 'smart' methods which use a more elaborate protocol between tag and reader. These more advanced tags will show more consistent behavior as the number of tags in the read zone increase.

The method of testing proved very practical, and it has the potential to be used for many more elaborate tests which were beyond the scope of this research. A natural next step would be to look at the reader's reaction to moving tags. Theoretically these read rates should be related to the sum of the static fields from the previous data for each path they follow, but it could well be that the results are different. It is generally accepted that moving tags are read better than static ones, and while one theory would say that is because they move between the zones plotted in Chapter 5, it is also conceivable that moving tags behave in a completely different manner.

What this research builds toward is a model for when and how a reader will read a tag, and the even more useful inverse: given a tag read, where is the tag most likely to be? Ideally this model could predict to high accuracy whether a read would occur for a given arrangement of reader, tag, and environment. Creating an ideal model is never possible, but it should not be impossible to develop a working model for RFID, the benefits of which might include, for example, more efficiently arranging readers to monitor automated environments.

7.2 Conclusions

Problems of identity and location sensing will exist for as long as machines interact with physical objects. The technologies will change, but the way people think about these problems will probably do so much more slowly, if at all. It is most important that as industries adopt technologies like RFID, they do so with an open policy and a system which meets their current needs without locking out future innovation. The analysis in this thesis is a practical example of the kind of further work which could be undertaken to improve a technology, even after it has been installed.

At present, RFID is an attractive and popular technology for automated environments. In many cases it is exactly what is needed. It provides adequate identity and location data for many applications. In some instances, however, it is not enough; it must be augmented by physical sensors or a higher-level algorithm to make its data useful to its employer. There are also situations where RFID is inadequate, and another technology must be used. With proper planning and installation RFID and other technologies can coexist within a single system, without having to make allowances for each other.

Appendix A

Raw Data

All of the data in this thesis were collected from the data logger into Zterm, a terminal emulation package. They were then processed using Perl scripts to strip characters and convert them to tensor form. The analysis was performed in Mathematica and made extensive use of Jens-Peer Kuska's excellent **MathGL3d** visualization tool, available from `http://phong.informatik.uni-leipzig.de/~kuska/mathgl3dv3/`.

A.1 Initial Data

All data are the number of reads in a five second interval. The first set of data was taken at 2cm intervals to provide a baseline with which to compare all subsequent datasets. They are plotted as a series of horizontal slices through the read zone, from the bottom to the top of the read zone.

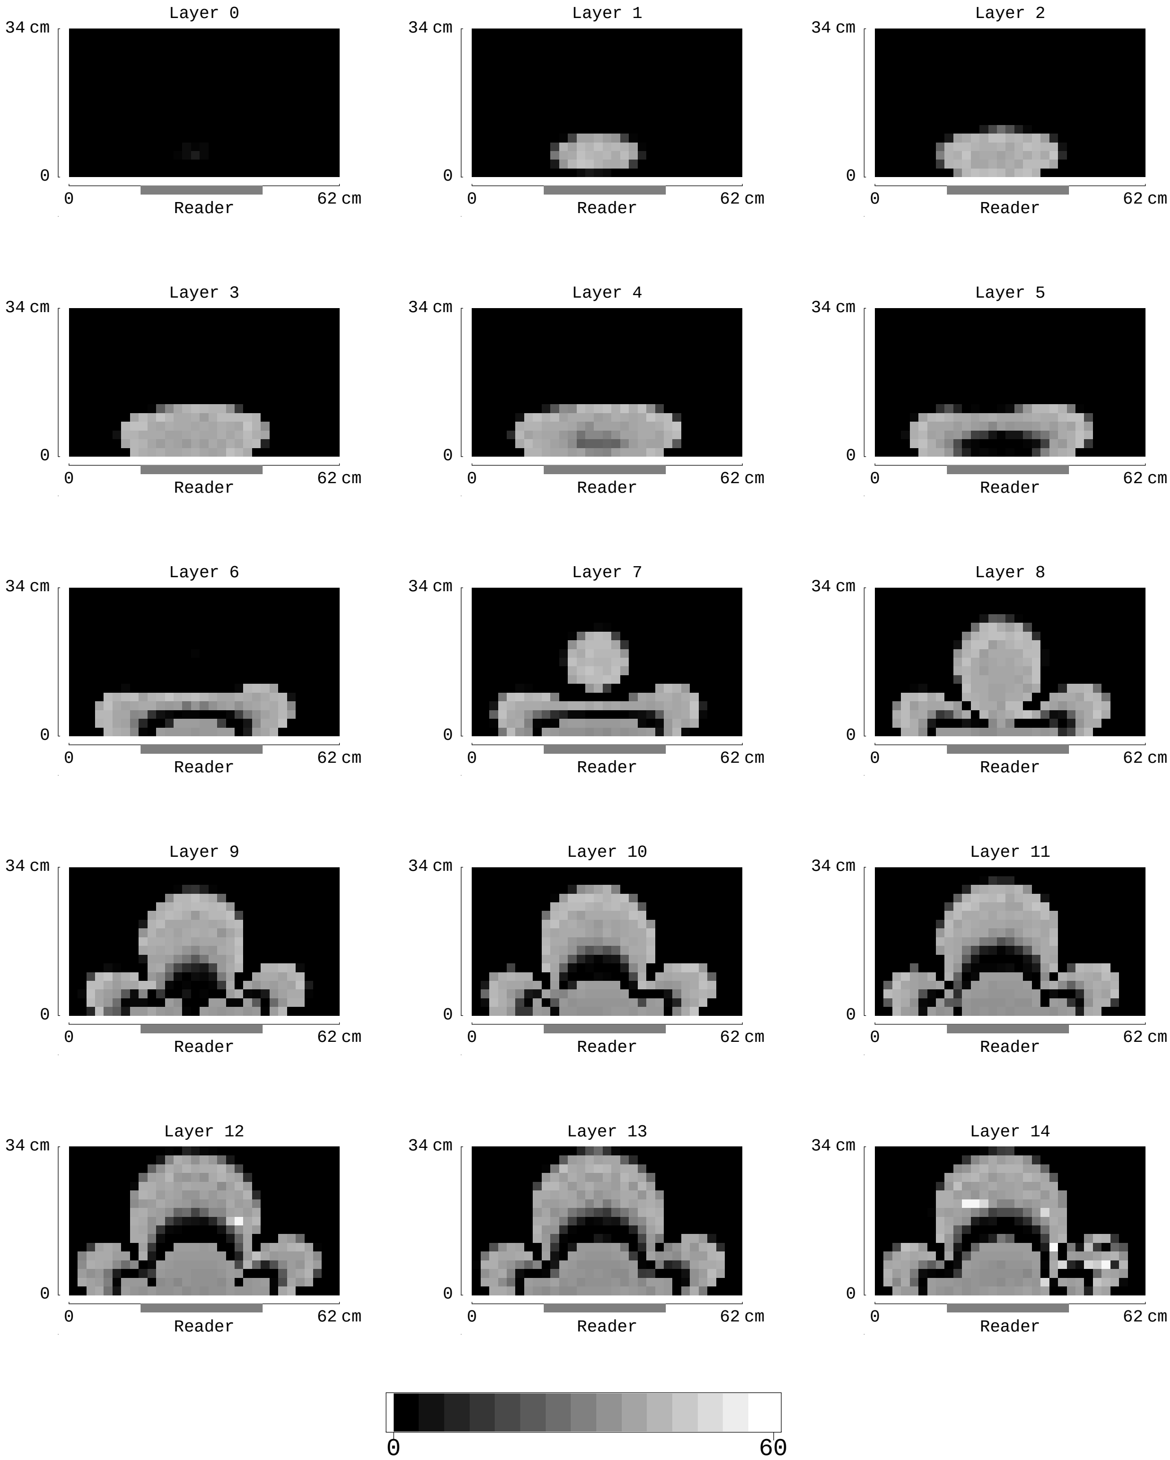

Figure A.1: Layers 0-14

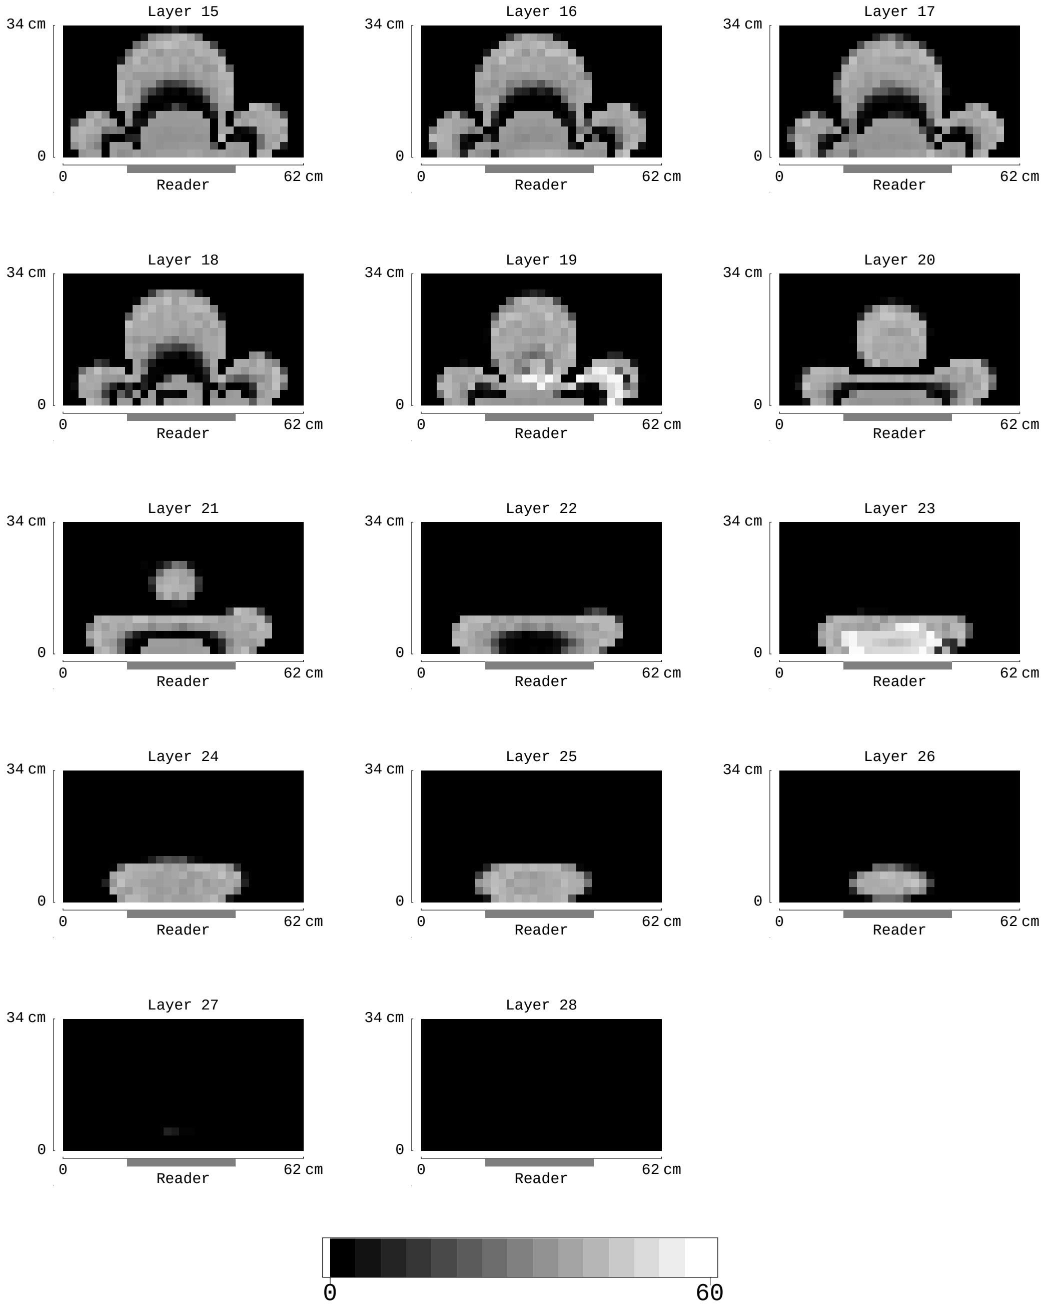

Figure A.2: Layers 15-28

A.2 Repeatability: 100 Trials

9 points were chosen in a large 3x3 grid in a region where they would provide a range of read rates. Each point was taken 100 times and the frequencies plotted. They have also been fitted with Normal distributions.

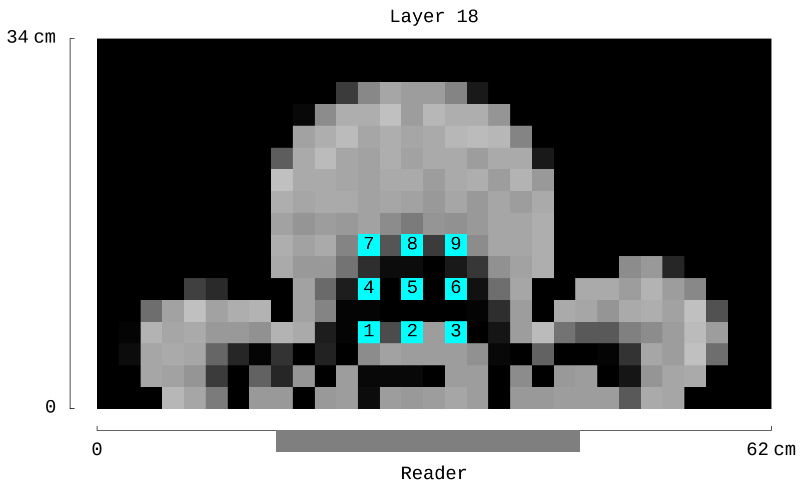

Figure A.3: The location of the 9 points

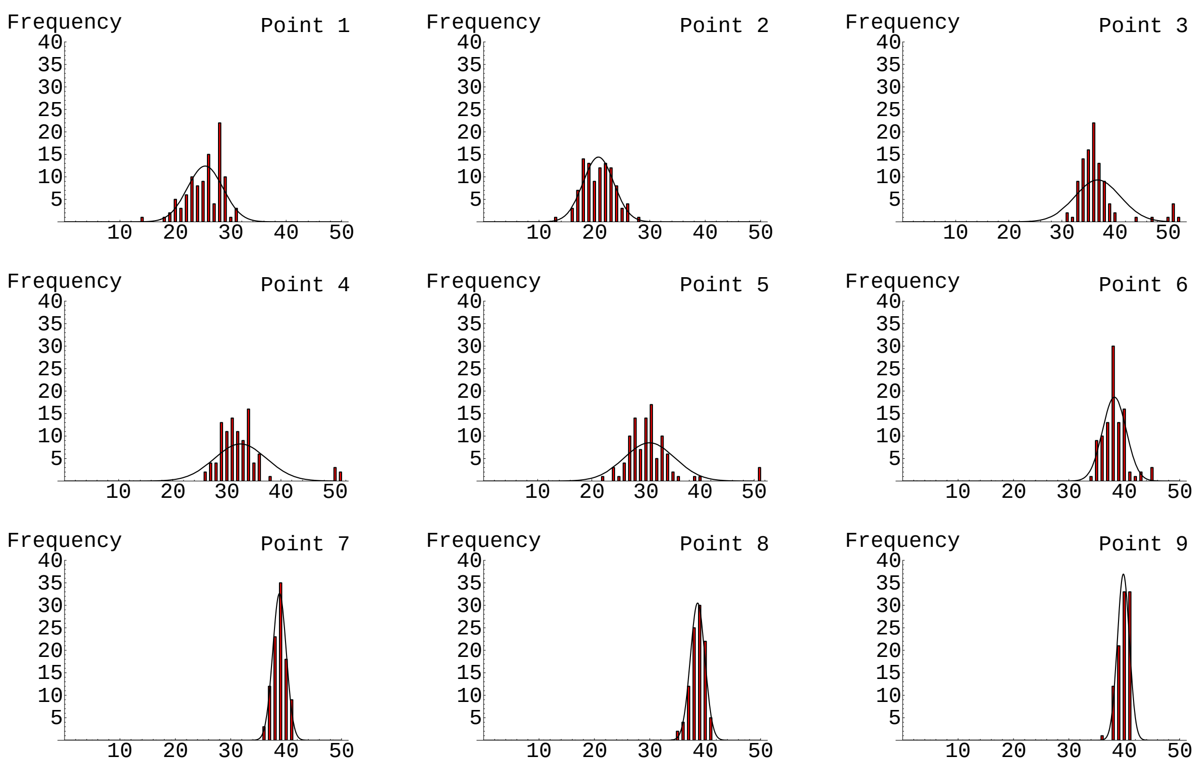

Figure A.4: The distribution of numbers of reads for 100 trials with 9 different points

A.3 Second set

This set was taken under the same conditions as the first, but data were taken at 4cm intervals in order to speed up the process to reduce the variation in the environment and collect them in such a time-frame that further data collection was feasible.

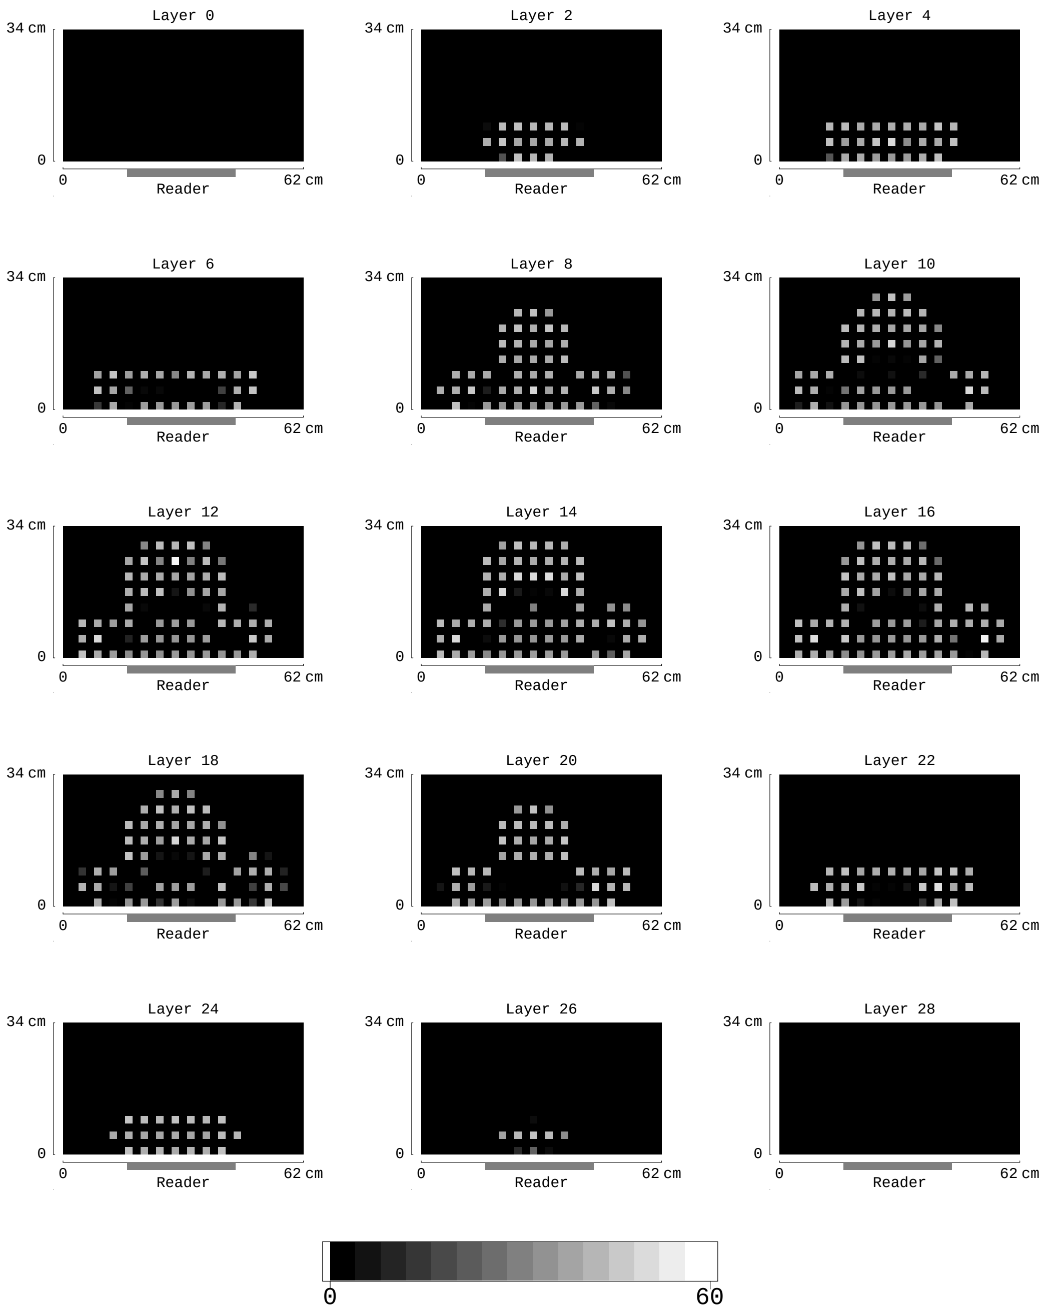

Figure A.5: Data at 4cm centers, identical configuration with the initial set

A.4 Long end effector

This set was taken with a long end effector installed on the robot to compare the results with the shorter one used in previous experiments.

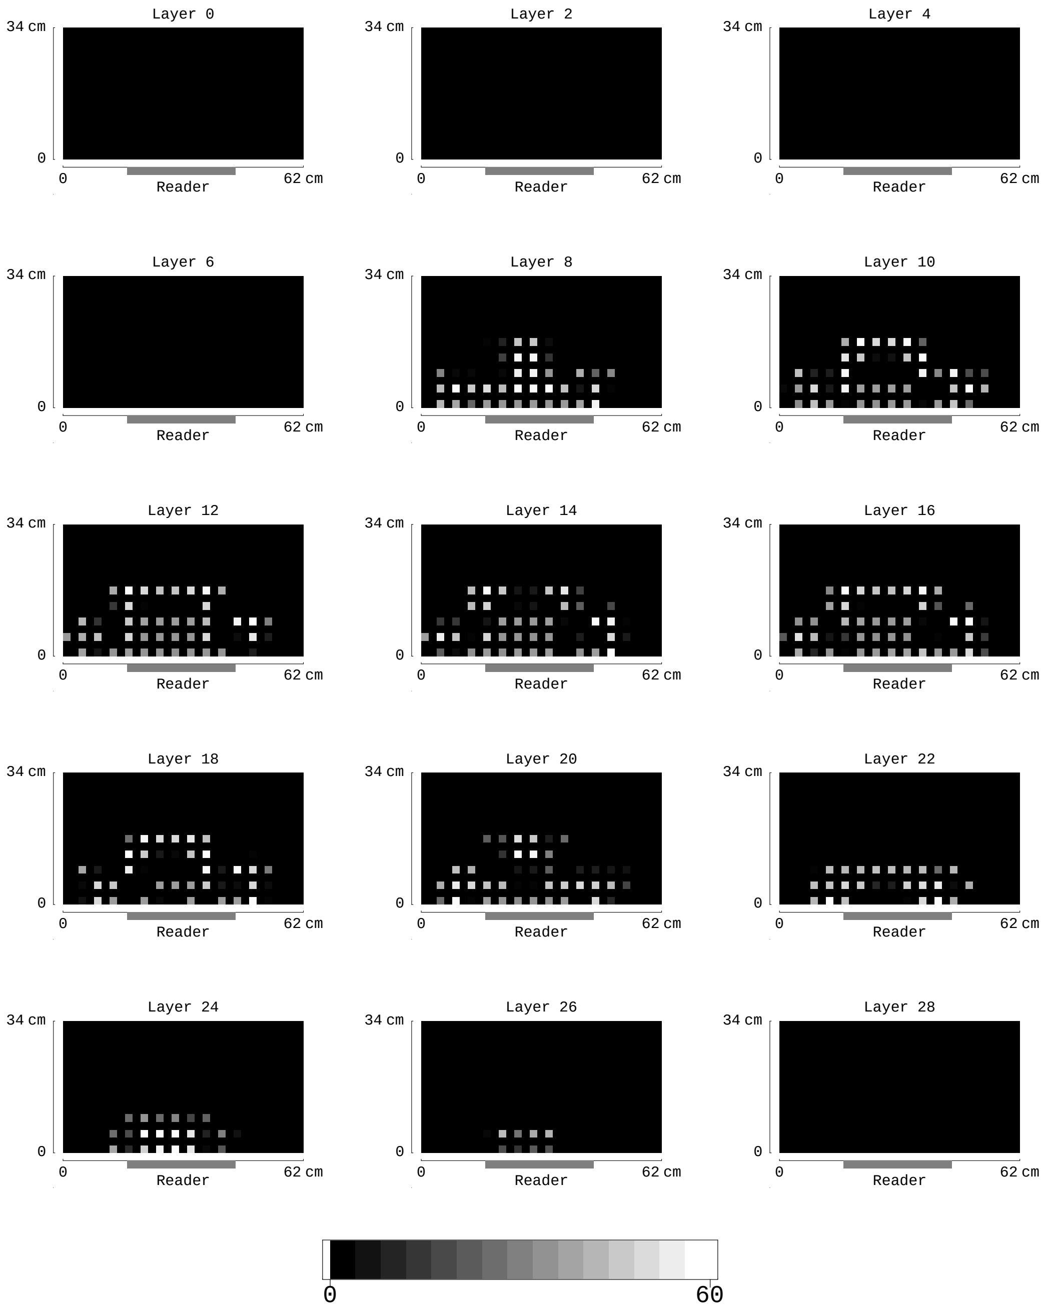

Figure A.6: Data at 4cm intervals with the longer end effector

A.5 Angled Tag

The tag was rotated 45° about the vertical axis to take this dataset. The angle made it impossible to collect data in the plane nearest the reader, which is why the graphs are all missing that line of data.

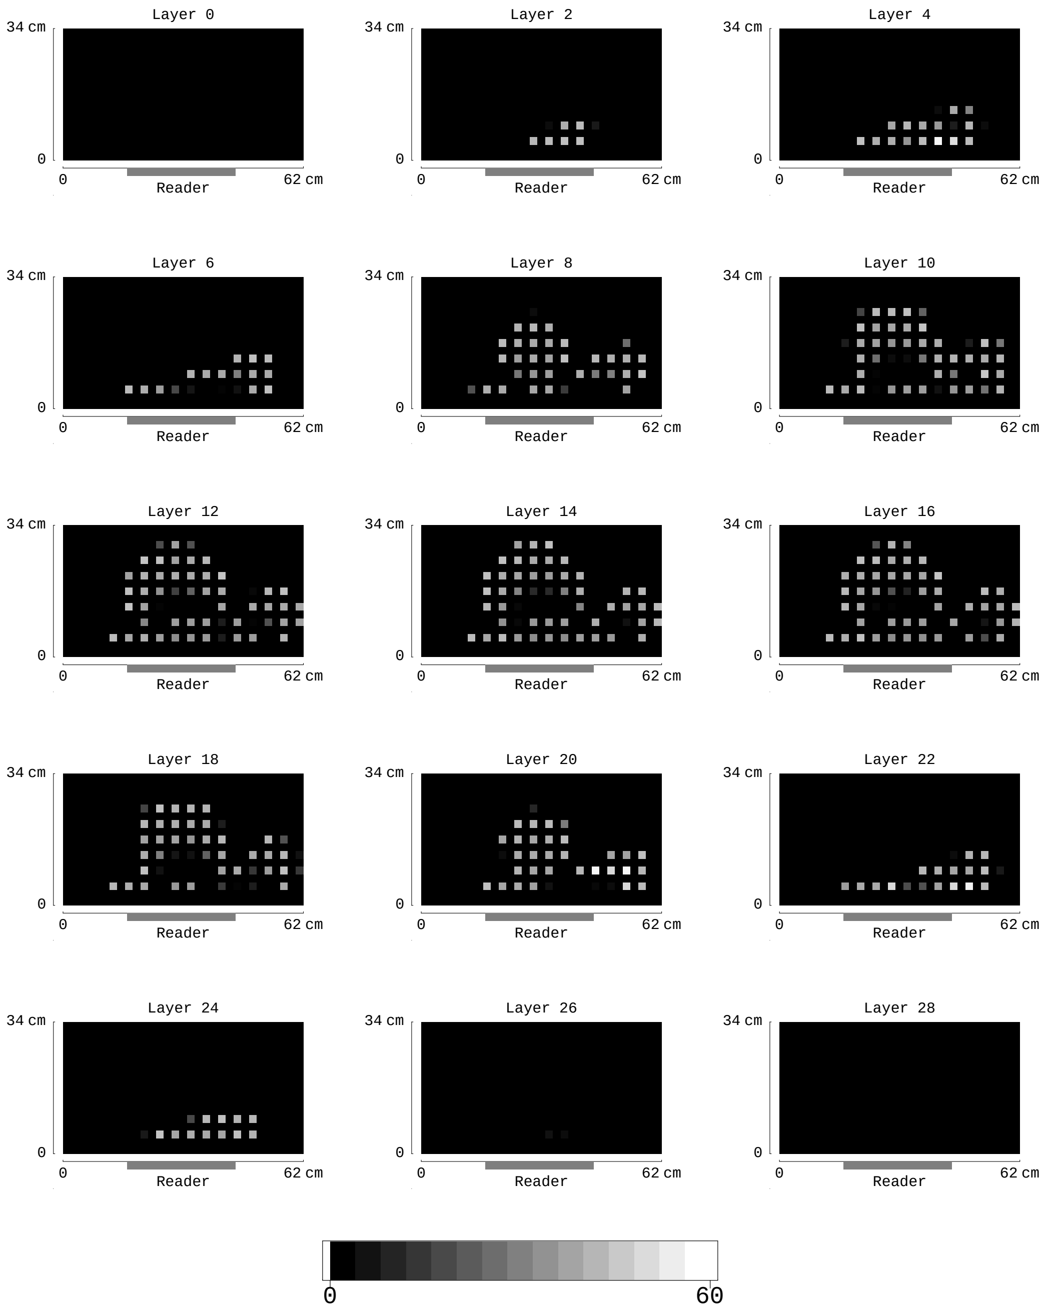

Figure A.7: Data at 4cm intervals with the tag at 45° to the reader

A.6 Large Tag

Checkpoint manufactures three sizes of tag: 5.3x5.3cm, 4.5x7.6cm, and 10x10cm. The largest size was used here, roughly 4 times the area of the tag used in the other experiments.

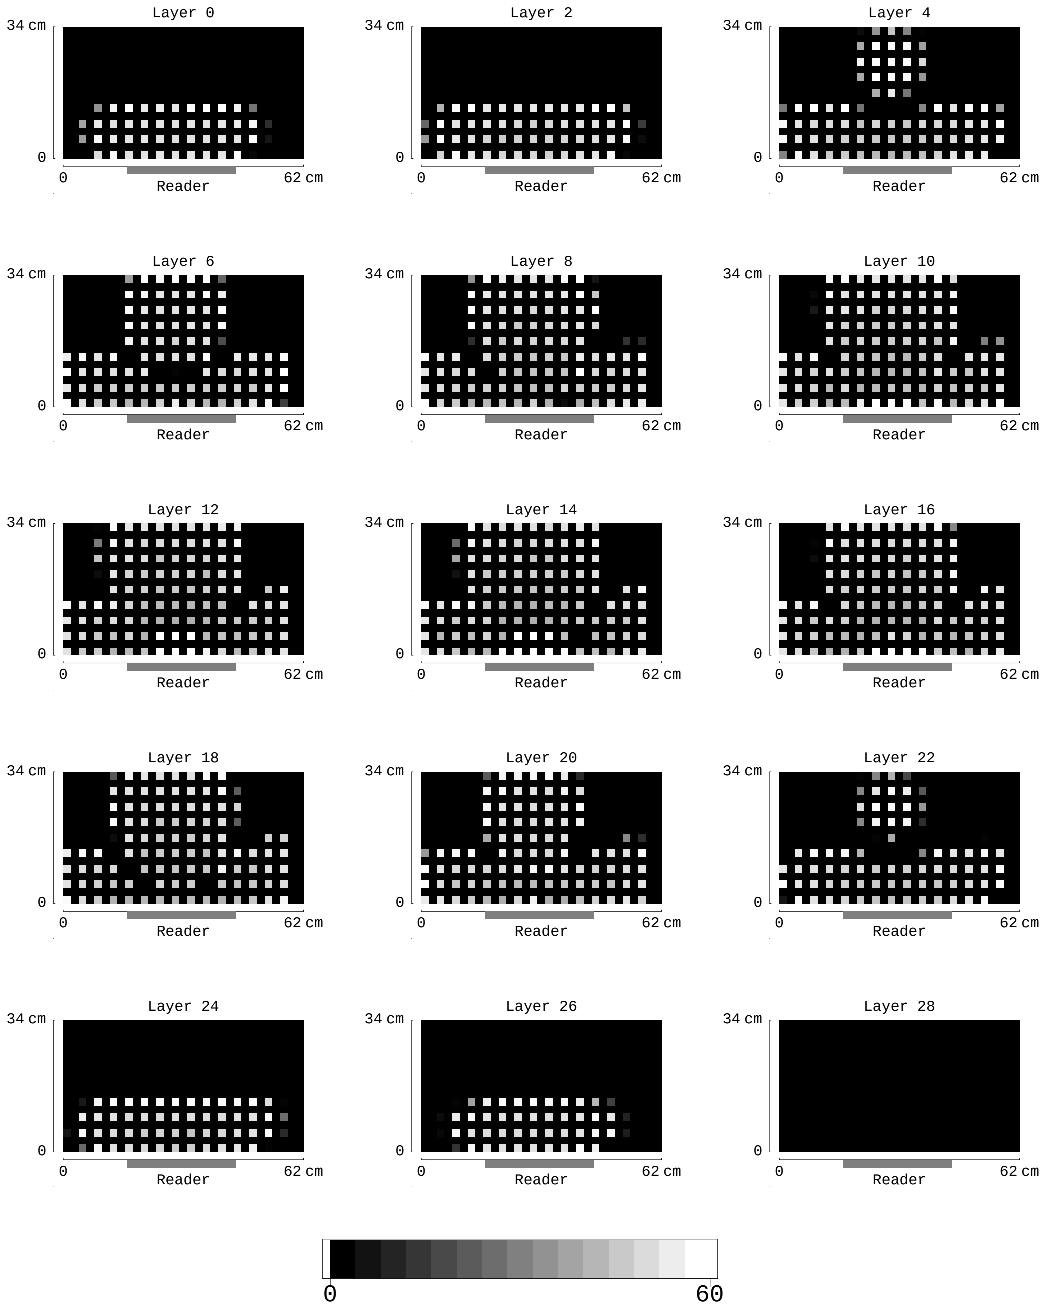

Figure A.8: Data at 4cm intervals with a 4" square tag

A.7 Collisions

This set had a second tag in read range to force the reader to do extra processing.

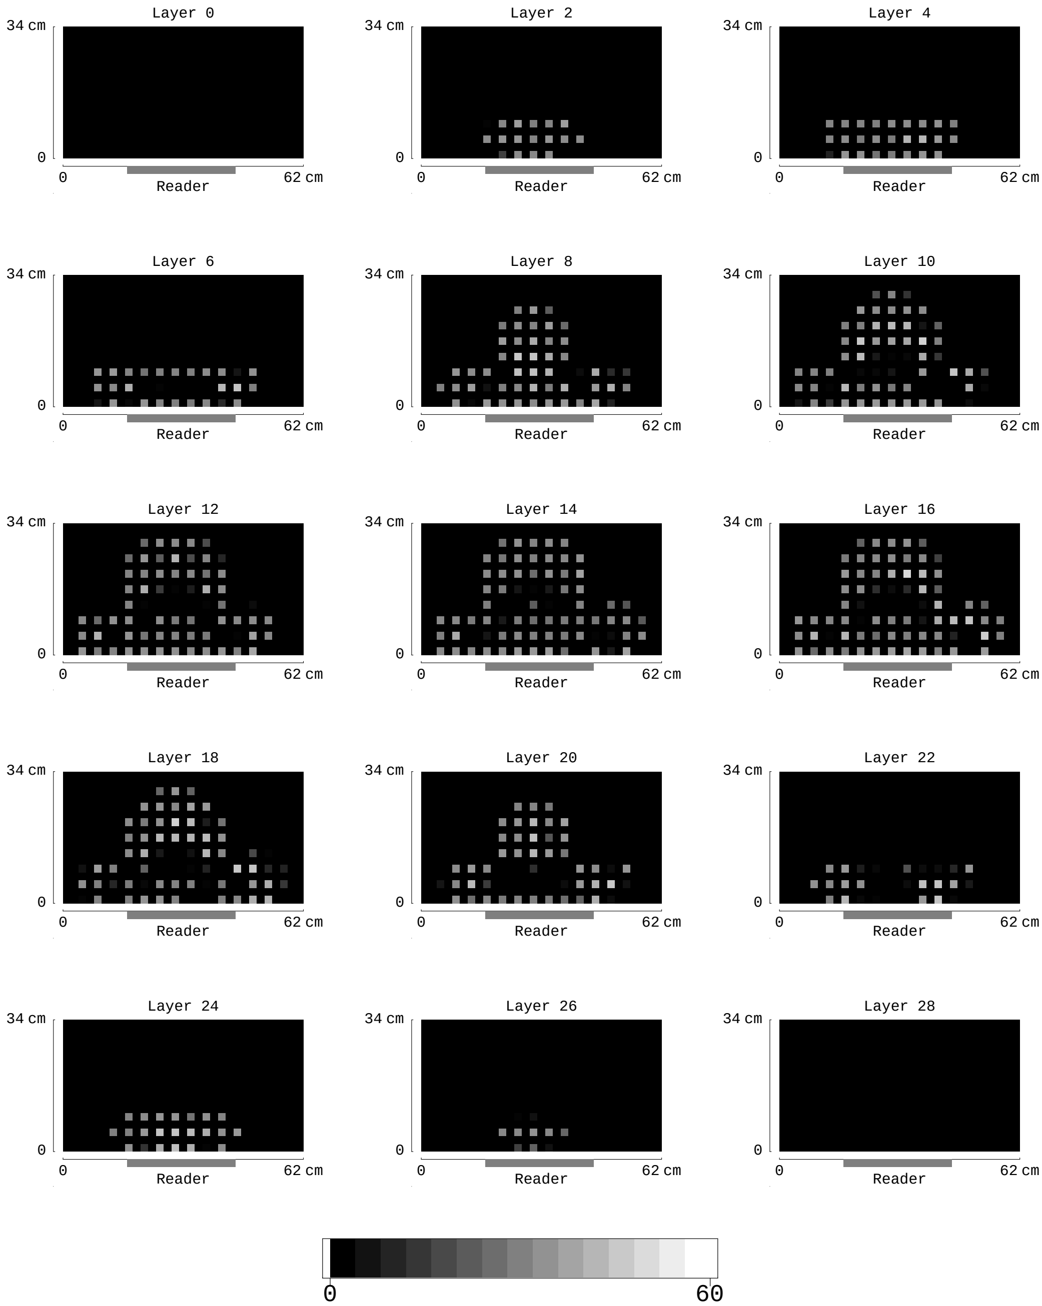

Figure A.9: Data at 4cm intervals with a second tag fixed in the read zone

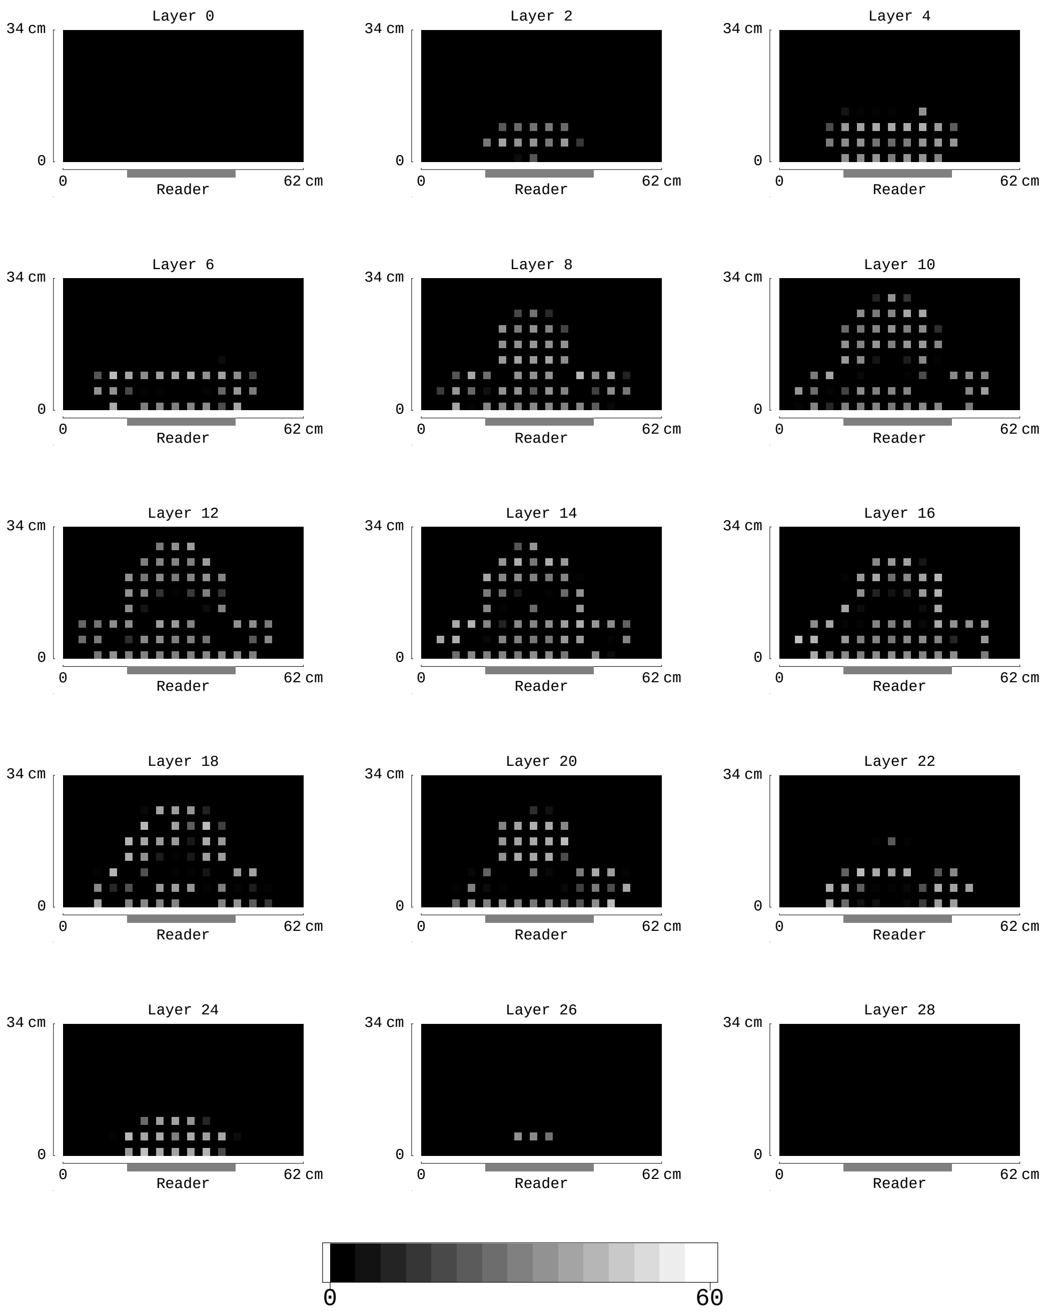

Figure A.10: Data at 4cm intervals with a second tag directly above the measured tag

A.8 Obstructions

Obstructions in the form of a steel bar were placed in front of the reader. The bar is shown in these renderings as a solid white line. It is only shown in the cross sections in which it appears, to make its position clear to the reader. In both sets of data, points where not collected in the area between the bar and reader at any elevation.

A.8.1 Bar flush with reader

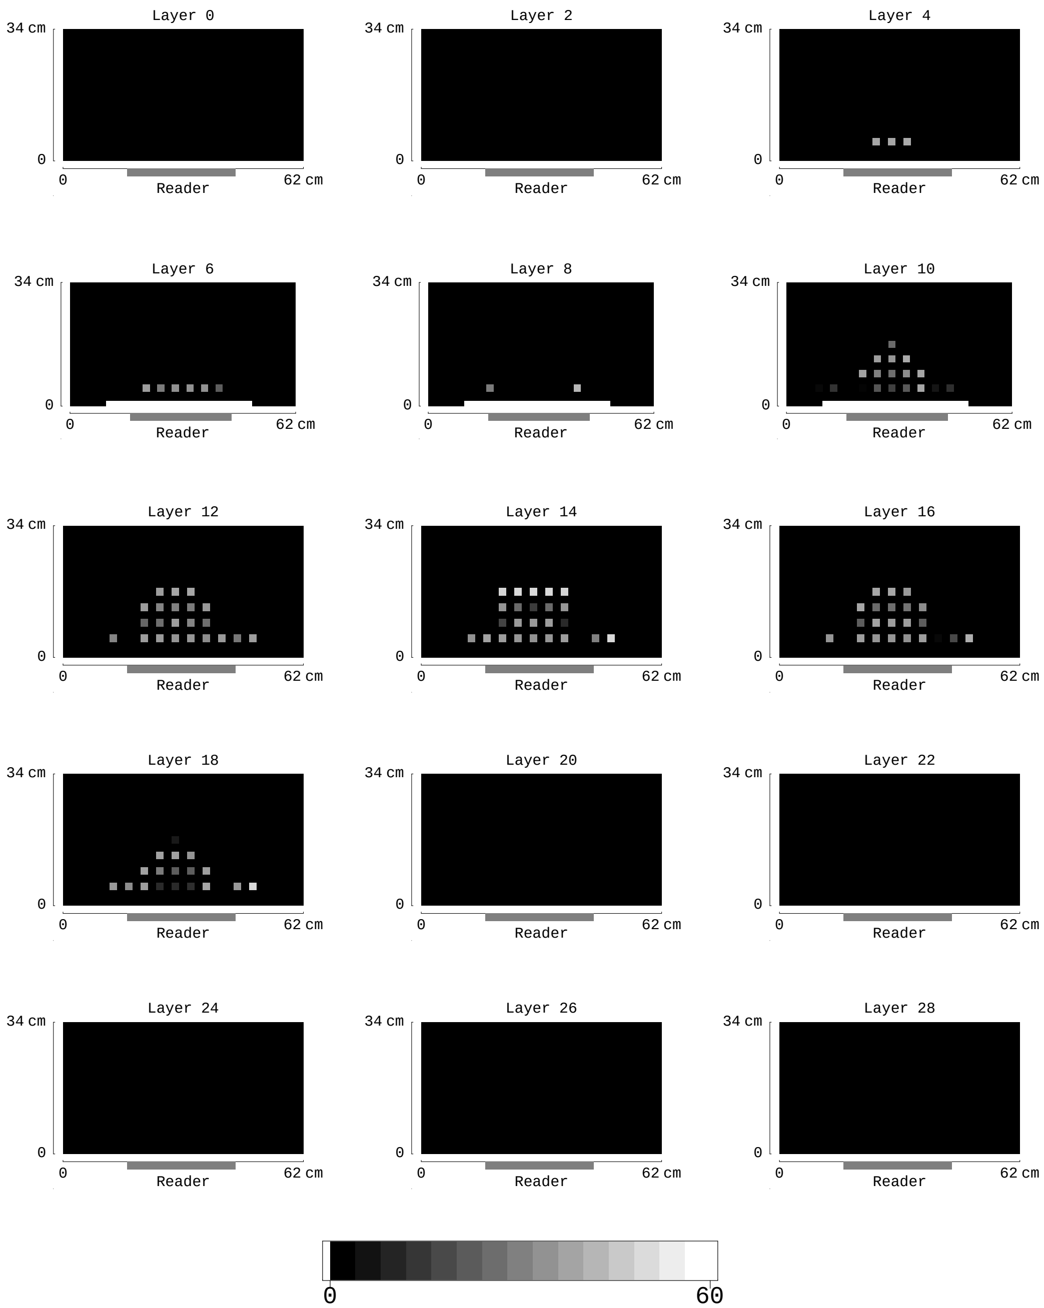

Figure A.11: Data at 4cm intervals with an 8x1.5cm bar flush with the bottom of the reader

A.8.2 Bar 10cm from reader

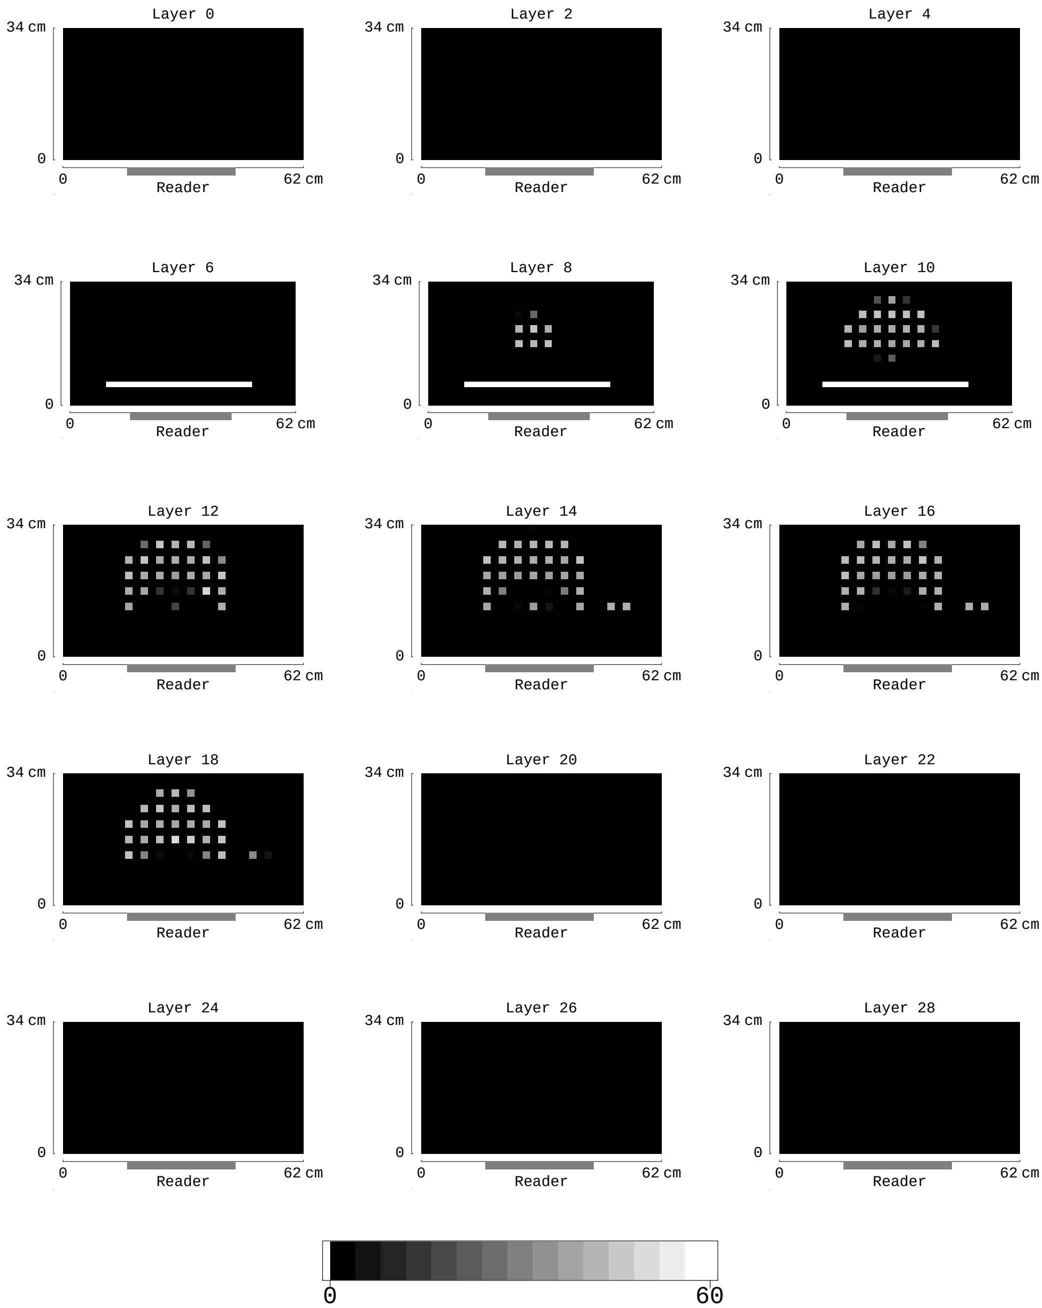

Figure A.12: Data at 4cm intervals with an 8x1.5cm bar 10cm from the reader

Appendix B

Source Listings

These Perl scripts were used to generate the program files for the robot and then process the data received back from the data logger. The script which performed the post-processing, `processdata.pl`, was written to be called from within Mathematica and uses Ulrich Pfeifer's Math::ematica module[1] to bridge between Mathematica and Perl. The first program, `rfid.pl`, generates the robot data files, which are then compiled with the Fanuc DOS compilation program `maketp.exe`.

Listing B.1: rfid.pl

```
#!/usr/bin/perl
#
# rfid.pl This program generates rfid?.ls, a series of Fanuc program
# files which can be compiled and run on the M6i. It takes a starting
# point, end point, step size in each dimension, tag orientation,
# robot arm speed, and the number of program files to divide the
# source file into.
# outputs: output.log summarizes what has happened
#          rfid?.ls are the program files
#          rfid?.ref map the point numbers from the program files
#          to world coordinates
if (($#ARGV + 1) < 17)  { ## not called from frontend

$xstepsize = 40;
$ystepsize = 40;
$zstepsize = 40;
$xmin = -859;
$xmax = -356;
$ymin = 795;
$ymax = 435;
$zmin = 135;
$zmax = 668;
$xoffset = 0;
$yoffset = 0;
$zoffset = 0;
$num_chunks = 2;
$lin_speed = 2000;
$w = -90;
$p = -86;
$r = -45;
}
else  {
$xstepsize = $ARGV[0];
$ystepsize = $ARGV[1];
$zstepsize = $ARGV[2];
$xmin = $ARGV[3];
$ymin = $ARGV[4];
$zmin = $ARGV[5];
$xmax = $ARGV[6];
$ymax = $ARGV[7];
$zmax = $ARGV[8];
$xoffset = $ARGV[9];
$yoffset = $ARGV[10];
```

[1] http://search.cpan.org/author/ULPFR/Math-ematica-1.108/

```
$zoffset = $ARGV[11];
$num_chunks = $ARGV[12];
$lin_speed = $ARGV[13];
$w = $ARGV[14];
$p = $ARGV[15];
$r = $ARGV[16];
}

$xsign = $xmax - $xmin;
$xsize = abs($xsign);
if($xsize == 0) {
  $xsign = 1;
}
else {
  $xsign = $xsign / $xsize;
}

$ysign = $ymax - $ymin;
$ysize = abs($ysign);
if($ysize == 0) {
  $ysign = 1;
}
else {
  $ysign = $ysign / $ysize;
}

$zsign = $zmax - $zmin;
$zsize = abs($zsign);
if($zsize == 0) {
  $zsign = 1;
}
else {
  $zsign = $zsign / $zsize;
}

$xstep = $xsize / $xstepsize;
$ystep = $ysize / $ystepsize;
$zstep = $zsize / $zstepsize;



if($xstep != int($xstep)) {
  $xstep = int($xstep) + 1;
}
if($ystep != int($ystep)) {
  $ystep = int($ystep) + 1;
}

if($zstep != int($zstep)) {
  $zstep = int($zstep) + 1;
}

$xstep++;
$ystep++;
```

```
$zstep++;


## Divide the processing into chunks

$chunk_size = $zstep / $num_chunks;
$leftover = 0;

if($chunk_size != int($chunk_size)) { ## catch the last little bit
  $leftover = $zstep - (int($chunk_size) * $num_chunks);
}

$chunk_size = int($chunk_size);

## make log file
$file = "./output.log";
open(INFO, ">$file");              # Open the file
$now = localtime;
print INFO "Fanuc␣robot␣program␣file␣generated␣$now\n";
print INFO "Start␣Point:␣(",$xmin+$xoffset,",",$ymin+$yoffset,",",
  $zmin+$zoffset,")\n";
print INFO "␣␣End␣Point:␣($xmax,$ymax,$zmax)\n";
print INFO "␣␣Step␣Size:␣x:$xstepsize␣mm␣y:$ystepsize␣mm␣z:$zstepsize␣mm\n";
print INFO "Datapoints␣in␣x:␣$xstep␣y:␣$ystep␣z:$zstep\n";

$cur_chunk_size = $chunk_size;

for($chunk = 1; $chunk <= $num_chunks; $chunk++)  {
    $zstart = ($chunk - 1) * $chunk_size;
    $zend = $zstart + $cur_chunk_size - 1;

    print INFO "Chunk␣$chunk:\n";
    print INFO "\tStart:␣(",$xmin+$xoffset,",",$ymin+$yoffset,",",
  $zmin+$zoffset + $zsign * $zstart * $zstepsize ,")\n";
    print INFO "\t␣␣End:␣(",$xmax+$xoffset,",",$ymax+$yoffset,",",
  $zmin+$zoffset + $zsign * $zend * $zstepsize,")\n";

    $count = 1;   # reset number of points, indexed for robot which starts at 1

    $lines = ($xstep * $ystep * $cur_chunk_size) * 4 + 1;

    $file = "./rfid$chunk.ls";
    $reference = "./rfid$chunk.ref";
    open(REF, ">$reference");
    open(FILE, ">$file");

    print REF "Position␣#␣:␣Matrix␣Element␣:␣World␣Position\n";

    print FILE "/PROG␣␣RFID$chunk\n/ATTR\nOWNER␣␣␣=␣MNEDITOR;\n" .
  "COMMENT␣␣␣␣=␣\"\";\nPROG_SIZE␣␣=␣3322;\n" .
  "CREATE␣␣␣=␣DATE␣03-05-06␣␣TIME␣14:07:00;\n" .
  "MODIFIED␣=␣DATE␣03-05-14␣␣TIME␣16:49:23;\n" .
  "FILE_NAME␣␣=␣;\nVERSION␣␣␣␣␣␣=␣0;\n" .
  "LINE_COUNT␣=␣$lines;\nMEMORY_SIZE␣␣=␣3606;\n" .
```

```
    "PROTECT␣␣␣␣=␣READ_WRITE;\nTCD:␣␣STACK_SIZE␣=␣0,\n" .
    "␣␣␣␣␣␣TASK_PRIORITY␣␣=␣50,\n␣␣␣␣␣␣TIME_SLICE␣=␣0,\n" .
    "␣␣␣␣␣␣BUSY_LAMP_OFF␣␣=␣0,\n␣␣␣␣␣␣ABORT_REQUEST␣=␣0," .
    "\n␣␣␣␣␣␣PAUSE_REQUEST␣␣=␣0;\nDEFAULT_GROUP␣=␣1,*,*,*,*;\n" .
    "CONTROL_CODE␣=␣00000000␣00000000;\n/MN\n";
      print FILE "␣␣␣1:␣␣R[200]␣=␣5␣;\n";
      for($i = 2; $i < $lines; $i +=4)  {
        printf FILE "%4d:L␣P[%d]␣%dmm/sec␣FINE␣␣;\n", $i,$count,$lin_speed;
        printf FILE "%4d:␣␣DO[7]␣=␣ON␣;\n",$i+1;
        printf FILE "%4d:␣␣WAIT␣R[200]␣;\n",$i+2;
        printf FILE "%4d:␣␣DO[7]␣=␣OFF␣;\n",$i+3;
        $count++;
      }

      print FILE "/POS\n";
      $count = 1;
      for($k = $zstart; $k <= $zend; $k++)  {
        for($j = 0; $j < $ystep; $j++)  {
          for ($i = 0; $i < $xstep; $i++) {
            print FILE "P[$count]{\n␣␣␣GP1:\n";
            print FILE "␣␣␣UF␣:␣0,␣UT␣:␣1,␣\tCONFIG␣:␣'F␣U␣T,␣0,␣0,␣0',\n";
            printf FILE "␣␣␣X␣=␣%5.2f␣mm,\tY␣=␣%5.2f␣mm,\tZ␣=␣%5.2f␣mm,\n",
              $xmin+$xoffset+$i*$xstepsize*$xsign,
              $ymin+$yoffset+$ystepsize*$j*$ysign,
              $zmin+$zoffset+$zstepsize*$k*$zsign;
            printf FILE "␣␣␣W␣=␣␣␣%2.2f␣deg,\tP␣=␣␣%2.2f␣deg,\tR␣=␣␣␣%2.2f␣deg\n};\n"
              $w,$p,$r;
            print REF $count-1,":($i,$j,$k)␣=>␣(",
              $xmin+$xoffset+$i*$xstepsize*$xsign,
              ",", $ymin+$yoffset+$ystepsize*$j*$ysign,",",
              $zmin+$zoffset+$zstepsize*$k*$zsign,")\n";
            $count++;
          }
        }
      }
      print FILE "/END\n";
      close(FILE);
      close(REF);
      if(($chunk == $num_chunks) && $leftover)  {
        $num_chunks++;
        if ($leftover > $chunk_size)  {
          $cur_chunk_size = $chunk_size;
          $leftover -= $chunk_size;
        }
        else  {
          $cur_chunk_size = $leftover;
            $leftover = 0;
        }
      }
}
close(INFO);                         # Close the file
```

The next program, processdata.pl imports a transcript of the output from the datalogger and strips out the data. It then outputs Mathematica code to input the data.

Listing B.2: processdata.pl

```
#!/usr/bin/perl

use lib qw(blib/lib blib/arch);

use Math::ematica qw(:PACKET :TYPE :FUNC);

my $ml = new Math::ematica;

sub new_matrix {

my $path = @_[0];
if (substr($path,-1) ne "/")  {
  $path = $path . '/';
}
# this is the step size for the original matrix (2cm centers)
$stepsize = 20;
$output = "";

open(REF, "<" . $path . "rfid.ref");
open(DATA, "<" . $path . "data.txt");

my(@lines) = <DATA>;           # read file into list
my(@refs) = <REF>;



if ($#lines == 0) {    # mac file format
  @lines = split(/\r/,$lines[0]);
}
shift @refs;
# these are the world coordinates for the robot when the tag is at
# the matrix origin, i.e. accounting for angle and end effector differences
open(OUTPUT,  "<" . $path . "output.log");
@output = <OUTPUT>;
close(OUTPUT);
($tmp,$xstart,$ystart,$zstart) = split(/Start Point: \(|,|\)/,$output[1]);
$output[4] =~ s/.*x: (.*) y.*z:(.*)/$1,$2/;
($xsize,$ysize) = split(/,/,$output[4]);

# find start of file
while (substr($lines[0],0,8) ne "Position") {
  shift @lines;
}
for($i=0; $i<=$#lines; $i++)  {
  $lines[$i] =~ s/Position:(.*)Reads:(.*)/$1,$2/;
  @line = split(/,/,$lines[$i]);
  $refs[$i] =~ s/.*=> \((.*)\)/$1/;
  chomp $refs[$i];
  $lines[$i] = "$refs[$line[0]],$line[1]";
  chomp $lines[$i];
  @data = split(/,/,$lines[$i]);
```

```
  @line = (($data[0] - $xstart) / $stepsize + 1,-($data[1] - $ystart) /
    $stepsize + 1,($data[2] - $zstart) / $stepsize +1, $data[3]);
  $output = $output . "t\[\[$line[2],$line[1],$line[0]\]\]␣=$line[3];\n";
  $output = $output . "mask\[\[$line[2],$line[1],$line[0]\]\]␣=1;\n";
}
close(REF);
close(DATA);
$output = $output . "Print\[\"Starting␣Point:␣($xstart,␣$ystart,␣$zstart)\"\];\n";
$output .= "LayerSize␣=␣" . $xsize * $ysize . ";\n";
return($output);
}

$ml->register('ProcessData', \&new_matrix,'String');
$ml->main;
```

The script is called in Mathematica like this:

```
...
t=Array[0&,{29,21,31}];mask=Array[0&,{29,21,31}];
link = Install["processdata.pl"];
ProcessData[x_String]:=ExternalCall[link,CallPacket[0,{x}]];
ToExpression[ProcessData["/path/to/data"]];
...
```

Arrays t and $mask$ are empty tensors the size of the total dataset. They are allocated with all zeros. For each data point that processdata recovers from the log file, it inserts it into t and sets the corresponding bit in $mask$ to 1. In this way all of the data sets collected have a standard size, regardless of which points are actually collected. The tensor $mask$ records which points are actually used so that when calculations are performed on a data set the user can differentiate between a value of 0 and the absence of a recorded value. Maintaining a separate mask tensor makes it easy to combine two masks so that only points shared by two tensors can be worked with. To this end a Mathematica function was written, TensorAnd, which performs a Boolean AND on corresponding elements of two tensors:

```
TensorAnd[thetensor_, themask_] := (TensorDim = Dimensions[thetensor];
    MaskDim = Dimensions[themask];
    If[Dimensions[TensorDim][[1]] < 3 || Dimensions[MaskDim][[1]] < 3,
      Message[TensorMask::"Invalid Dimensions"];Abort[]];
    If[TensorDim != MaskDim, Message[TensorMask::"Dimension Mismatch"];Abort[]];
    NewTensor = Array[0 &, TensorDim];
    For[i = 1, i <= TensorDim[[1]], i++,
      For[j = 1, j <= TensorDim[[2]], j++,
        For[k = 1, k <= TensorDim[[3]], k++,
          If[(themask[[i, j, k]] == 1) && (thetensor[[i,j,k]] == 1),
            NewTensor[[i, j, k]] = 1]]]];
    Return[NewTensor];);
```

During the investigation it was desirable to ignore the large area of 0-valued points around the read rate data. To do so a mask had to be created which masked out the large areas of 0 readings around the data. There are potentially individual 0 readings within the data, however, which should not be masked out. A Mathematica function was written which created a mask which masked out all of the zeros which were surrounded on all sides by zeros. This gave very satisfactory results, successfully masking out the data which was clearly beyond the range of the reader, but leaving the individual 0s within the data field. The function used was as follows:

```
MaskZeros[thetensor_] := (TensorDim = Dimensions[thetensor];
 If[Dimensions[TensorDim][[1]] < 3,
   Message[MaskZeros::"Invalid Dimensions"];Abort[]];
   ReturnMask = Array[1 &, TensorDim];
    For[i = 1, i <= TensorDim[[1]], i++,
     For[j = 1, j <= TensorDim[[2]], j++,
      For[k = 1, k <= TensorDim[[3]], k++,
       lessx = False; morex = False;
       lessy = False; morey = False;
       lessz = False; morez = False;
       If[i == 1, lessx = True, If[thetensor[[i - 1, j, k]] == 0, lessx = True]];
        If[i == TensorDim[[1]], morex = True, If[thetensor[[i + 1, j, k]] == 0,
            morex = True]];
         If[j == 1, lessy = True, If[thetensor[[i, j - 1, k]] == 0, lessy = True]];
          If[j == TensorDim[[2]], morey = True, If[thetensor[[i, j + 1, k]] == 0,
              morey = True]];
           If[k == 1, lessz = True, If[thetensor[[i, j, k - 1]] == 0, lessz = True]];
            If[k == TensorDim[[3]], morez = True, If[thetensor[[i, j, k + 1]] == 0,
                morez = True]];
             If[lessx && morex && lessy && morey && lessz && morez,
                 ReturnMask[[i, j, k]] = 0];]]];
    Return[ReturnMask]);
```

The data came into the computer from the datalogger: a PIC 16F876 microcontroller which processed the serial output from the reader and produced the number of reads data for the PC. The logger was controlled by an I/O line from the robot controller. The program was written in CCS C.

Listing B.3: rfid.c

```
/* rfid.c -- logs data from a Checkpoint RFID reader
 * data comes in the format X<18-digit ID number>N
 * with no breaks.
 * This program simply counts the Ns in a stream of data
 * while a read-enable line is high. This line is controlled
 * by the robot controller, moves the RFID tag in place and then
 * raises the line for a specified period of time. The data is logged
 * to a PC.
 */

#include <16f876.h>
#fuses HS,NOBROWNOUT,NOWDT,NOPROTECT,NOLVP
#use delay(clock=20000000)
#use rs232(baud=38400,xmit=PIN_C6,rcv=PIN_C7,restart_wdt,ERRORS)

#byte RCREG = 0x1
#byte porta = 0x05
```

```
#byte portb = 0x06

void main(void)   {
   char character = 0, last_char, output = 0;
   int16 read_count = 0, tag1_count = 0, position = 0;
  int1 logging = 0;

  set_tris_b(0x01);   // use B0 for trigger input
  set_tris_c(0B10000000);
  bit_set(portb, 1);

  printf("Checkpoint␣RFID␣Data␣Logger\n\r");

   while(TRUE) {
    // these two empty the receive buffer if it's filled up
    // while B0 was low
    if(kbhit())
      getch();
    if(kbhit())
      getch();
    if(portb & 1)
      bit_set(portb, 1);   // light an LED when recording data
    if(portb & 1) {
      delay_ms(20);
      logging = 1;
    }
      while(logging)  // while we're taking data
      {
      if(kbhit()) {
        last_char = character;
           character = getc();
//         putc(character);
        delay_us(26); // simulate a putc, the loop seems wound too tight
                  // without this delay and performs erratically.
          if(character == 'N')  {
          read_count++;
          if(last_char == '0')
            tag1_count++;
        }
      }
      output = 1;   // Output when we're out of this loop
      if((portb & 1) == 0)  {
        delay_ms(10);
        logging = 0;
      }
      }
      bit_clear(portb, 1);    // clear LED
    if(output)  {
        printf("Position:␣%4lu␣Reads:␣%3lu␣Tag␣0:%3lu\n\r", position,
    read_count, tag1_count);
      read_count = 0;
      tag1_count = 0;
      position++;
      output = 0;
```

```
    }
   }
}
```

Bibliography

[1] RFID: Technology evolution, market segmentation analysis and player profiles. Technical report, Allied Business Intelligence, Inc, 2002.

[2] Mike Addlesee, Rupert Curwen, Steve Hodges, Joe Newman, Pete Steggles, Andy Ward, and Andy Hopper. Implementing a sentient computing system. *IEEE Computer*, 34(8):50–56, 2001.

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

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

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

[6] N. Bulusu, J. Heidemann, and D. Estrin. GPS-less low cost outdoor localization for very small devices, 2000.

[7] 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.

[8] Jeffrey Hightower and Gaetano Borriello. Location systems for ubiquitous computing. *IEEE Computer*, 34(8):57–66, August 2001.

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

[10] Todd D. Hodes, Randy H. Katz, Edouard Servan-Schreiber, and Lawrence A. Rowe. Composable ad-hoc mobile services for universal interaction. In *Mobile Computing and Networking*, pages 1–12, 1997.

[11] ID-TECK,Inc. Proximity systems frequently asked questions. ID-TECK customer support web site. `http://www.idteck.com/english/support/faq03.html#5`.

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

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

[14] Jeff Jacobsen. Quoted in RFID Journal. Technical report, Alien Technology, September 24, 2002. http://www.rfidjournal.com/article/articleview/76/1/1/.

[15] 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.

[16] Jeremey Landt. Shrouds of time: The history of rfid. Technical report, The Association for Automatic Identification and Data Capture Technologies, 2001.

[17] Microchip. *Microchip MCRF355 Datasheet*. http://www.microchip.com/download/lit/pline/rfid/devices/21287f.pdf.

[18] Christine Spivey Overby. Wal-mart drops item-level rfid pilot: Big whoop. Technical report, Forrester Research, 2003. http://www.forrester.com/ER/Research/Brief/0,1317,17140,00.html.

[19] 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.

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

[21] Fanuc Robotics. *Fanuc M6ib Datasheet*, 2003. http://www.fanucrobotics.com/datasheets/robots/M-6iB.pdf.

[22] 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.

[23] D. Simplot, M. Latteux, and R. Kalinowski. An adaptive anti-collision protocol for smart labels. Joint Gemplus and LIFL contribution to ISO/IETC 18000-3 Working Group, 2001.

[24] Checkpoint Systems. *Performa RFID Tag Datasheet*, 2003. http://www.checkpt.com/docs/rfidtag.pdf.

[25] Checkpoint Systems. *Performa Slimline RFID Reader Datasheet*, 2003. http://www.checkpt.com/docs/rfidslim.pdf.

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

[27] Texas Instruments. Zebra Technologies and Texas Instruments RFID Systems announce technology collaboration for secure 13.56MHz access control solutions, 2002. http://www.ti.com/tiris/docs/news/news_releases/2002/rel9-10-02b.shtml.

[28] Knowledge centers: Radio frequency identification. Technical report, Texas Instruments, 2002. http://www.ti.com/tiris/docs/news/in_the_news/2002/12-16-02.shtml.

[29] PulsON technology overview. Technical report, Time Domain Corporation, 2001. http://www.timedomain.com/Files/downloads/techpapers/PulsONOverview7_01.pdf.

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

[31] Trakus, inc., 2003. http://www.trakus.com/svc_hardware.html.

[32] Harald Vogt. Efficient Object Identification with Passive RFID Tags. In *International Conference on Pervasive Computing*, LNCS. Springer-Verlag, 2002.

[33] 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.

[34] Roy Want and Daniel M. Russell. Ubiquitous electronic tagging. *IEEE Distributed Systems Online*, 1(2), 2000.

[35] A. Ward, A. Jones, and A. Hopper. A new location technique for the active office, 1997.

[36] G. Welch, G. Bishop, L. Vicci, S. Brumback, K. Keller, and D. Colucci. High-performance wide-area optical tracking: The hiball tracking system, 2001.

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