Aero-ID Sensor Integration: Scope of Work

James Brusey, Alan Thorne
Auto-ID Lab, University of Cambridge, UK

This report presents the scope of work undertaken in the RFID (Radio Frequency Identification) Sensor Integration theme of the Aero-ID Project. This theme arises from the realisation that establishing the identity of a product or component is often not enough. There is a growing awareness that to manage our resources efficiently and effectively, it is necessary to also know the condition or state of the component. On its own, RFID cannot help us and other types of sensor will be required. For many application areas, establishing identity will be essential to the process of converting sensory data into knowledge about the condition of a particular component. In other words, the result of "fusing" the data from regular and RFID sensors will be more informative than that obtained from either one alone.

In this report we examine this idea in more detail, provide reasons for why this research theme is worth studying, and propose a methodology for investigating it in the context of the aerospace industry. Central to this proposed methodology is interaction and engagement with the aerospace industry to ensure that the resulting research is relevant and couched in terms that will be understandable to that sector. To achieve this, the early work will define a set of characteristic axes to map out the different applications of RFID sensor integration and to attempt to discover common elements.

AUTO-ID LABS AEROID-CAM-003

Published February 1, 2006. Distribution restricted to sponsors until August 1, 2006

Contents

Published February 1, 2006. Distribution restricted to sponsors until August 1, 2006

1 Introduction

Radio Frequency Identification (RFID) is typically treated in isolation from all of the other things that may be going on in a system. As a sensory device, the data from it are generally not integrated with other data sources. RFID has mainly been used to establish an idea of the *location* of some set of physical objects. As far as the physical state of the object is concerned, it can really only estimate its location and not any other aspect of its current condition. For many application areas, knowledge of identity or location is insufficient; other types of sensors are required. Even though other sensors must be used, RFID remains relevant. In many cases, identity and location remain important and are often critical to correctly inferring the meaning of data from other sensors.

1.1 Motivation

RFID is a rapidly maturing area. Standards are being developed, and in a field that had previously seen the proliferation of customised and individual solutions, we are now seeing the rise of a unified approach based around EPCglobal's EPC Network [12]. Nevertheless, one size does not fit all. In particular, the EPC Network does not currently address integrating sensors with RFID. Integrating sensor data is critical to correctly dealing with mobile items whose physical state is important. RFID on its own can tell us about location or identity, but not state. Other types of sensors, however, may be unable to detect when components are changed over, or an old product is reused. For these reasons, the lack of developed standards and approaches to RFID sensor integration *forms a barrier to wider adoption* of RFID technologies.

We believe that the key question from the Aero-ID sponsors' perspectives is *not* whether to use RFID and sensors together, as for many applications, this will tend to happen anyway. Rather the question is one of *how to do so to our best advantage*; to minimise costs, while maximising the effectiveness and quality of the information derived.

For this research theme to properly answer this question, it must face a number of challenges. First, whatever tools, methods, or approaches are developed, they must be *sufficiently general* so that they may be used for a wide range of applications. Specifically, it should make sense to evolve the developed approaches into standards. Second, the developed approach must be *practical*. In other words, as far as possible within the constraints of available resources, the approach should be made to be useful and usable in an industrial setting. Finally, the approach should be made to provide results that are close to the *best possible*. That is, it is desirable to gain a large, and hopefully the largest possible, improvement by applying the developed approach.

In order to perform this research, we propose a series of steps as outlined in Section 4. We expect that the output of this research program will include a well documented approach to RFID sensor integration, along with laboratory and industrial demonstrations. It is planned to draw on the experience of a number of industry partners in the Aero-ID consortium, as well as expertise from the University of Cambridge, and other Auto-ID Labs.

1.2 Report Overview

Figure 1 provides a schematic of the report structure. The next section looks at industry applications that have either already integrated RFID and other sensor data or where such integration seems likely to be useful. Section 3 provides background to the main issues affecting this work. Section 4 proposes an overall methodology to be followed by this research theme. Section 5 concludes the report.

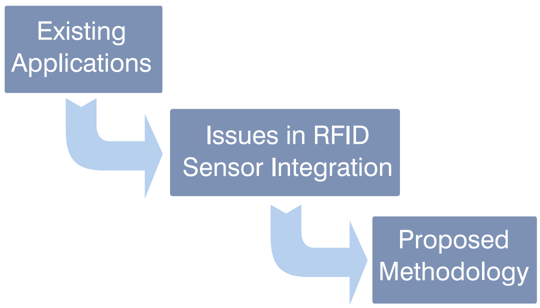

Figure 1: Structure of this Report

2 Industrial Applications of RFID Sensor Integration—A Review

What sorts of things might RFID sensor integration be used to do? Some candidate application areas identified so far in the Aero-ID program include monitoring shipping containers, and helping with Integrated Vehicle Health Management (IVHM). These two applications are now discussed in more detail.

Published February 1, 2006. Distribution restricted to sponsors until August 1, 2006

2.1 Container Security Initiative

The US Customs Service has started a Container Security Initiative designed to prevent the smuggling of terrorist weapons in cargo containers. One part of this initiative involves the use of tamper-detection seals on high-risk containers. As Stanford [11] points out, active RFID tags might be used to monitor the electronic seals during transit, as well as monitor other environmental factors such as temperature and humidity. Although Stanford mentions that the RFID tag might conceivably transmit alerts to shippers or customs authorities, this link would likely be handled via satellite. The main use of RFID would rather be at either end of the transit, where customs officials would interrogate the tag to ensure that no tampering had occurred, and the shipper may need to check that the environmental conditions during transit were within limits.

2.2 Integrated Vehicle Health Management

Integrated Vehicle Health Management (IVHM) is to aerospace what on-board diagnosis (OBD) are to the automotive industry. Ofsthun [7] points out that there are four primary areas that IVHM can address: diagnosis, prognosis, automated inspections, and anomaly detection (see Section 3.3 on page 12 for further discussion of these areas). It is a little unclear as yet what role RFID has to play in IVHM, however one possible application is to correctly handle the replacement of component parts of the vehicle. In principle, collected data for a component part should travel with that part. However traditional on-board diagnosis approaches tend to centralise data collection in a single embedded computer. If we assume that all collected data is transmitted to a back-end database, it is conceivable that the profile of a single component can be correctly tracked by using some form of unique identifier for that component. RFID could be used to provide that identifier. An alternative approach is for each component to collect data locally on the RFID tag. IVHM functionality may still require some transfer of that data to a central computer but, by keeping a local copy of the data, it is then simpler to ensure that logged data travels with the component.

Using RFID sensor integration in IVHM may be particularly relevant for structural monitoring in hard-to-reach locations (e.g. corrosion under galleys or floor panels). Although it is conceivable that such sensors might be connected to the network, it may be difficult in practise to achieve this. The use of small RFID sensor packages instead of network connected sensors may avoid costly redesign to route network cables, and provide safety certification.

2.3 Other Applications

There are a number of existing applications that combine RFID and sensors. The following is not intended to be an exhaustive list but representative of the variety of the applications. In comparison with the previously mentioned applications, these have not yet been identified by consortium members as being key.

2.3.1 RFID and Temperature Sensing of Perishables

A number of companies, such as InfraTab [2] and CliniSense [1, 3], are focusing on the problem of using RFID-based temperature sensors to answer questions about the freshness and quality of perishable goods. Their insight is to combine temperature and time duration (literally *integrating* temperature over time) to derive an estimate of bacterial growth, and thus derive an accurate estimate of the freshness of the product. As Collins [3] points out, a key advantage of RFID sensors is that they are inexpensive, and can provide more accurate information than many other approaches.

2.3.2 EU PROMISE Project

PROMISE [4] is an European Union funded project involving a consortia of 22 partners (including the University of Cambridge) over 9 countries. Its aim is to examine the role of smart embedded devices to monitor and track a product throughout its life-cycle. This project is expected to impact beginning, middle and end-of-life cycle operations by allowing information to flow about a product through the different stages. It involves 11 industrial demonstrators, as well as a training demonstrator. The industrial demonstrators range from supporting predictive maintenance on gas boilers (c.f. IVHM prognosis), through "Design for X" for heavy transport vehicles, to intelligent reuse of automotive clutch pads. A central aspect of this work is to incorporate a Product Embedded Intelligent Device, which is effectively an RFID tag plus sensors and memory, and to show how this can be used to provide sufficient information to make real improvements in the way products are made, used, and recycled. Although none of the demonstrators are specifically aimed at the aerospace industry, we believe that the ideas and approaches will be highly applicable to it.

2.3.3 Logistics

For the most part, logistical requirements are adequately dealt with by using RFID alone. The main benefits of incorporating other sensors is to increase the accuracy with which a container or part can be tracked. For example, by integrating GPS data with RFID and fitting satellite or GPRS communication into the transporter or warehouse, it becomes possible to accurately track shipments while they are in transit. Another example uses time-of-flight information of radio signals in a WiFi network to triangulate and thus provide an estimate of the location of a tag within a few centimetres.

3 Issues in RFID Sensor Integration

In this section, we now turn to the fundamental problems in dealing with RFID and sensor integration.

3.1 Connecting Sensors to Tags

In order to combine data from RFID and other types of sensors, it is necessary that the two disparate types of technology be brought together. This can either be done in a *physical* sense, where we attach sensors directly to the RFID tag, or it may be that the sensors are not connected, but we may still need to develop approaches to allow the two to be *logically* unified.

3.1.1 Physically Connected Sensors

There are a number of existing technologies for connecting sensors to RFID tags. Generally speaking, active (battery powered) tags are used since the sensors require some power to operate and may need to be active when not within the read field of an RFID antenna. Figure 2 shows the generic components of a typical RFID sensor package. As with standard RFID tags, a basic RFID chip and an antenna is required. There may be additional logic and memory, a clock to record the time of events, one or more sensors, and a battery.

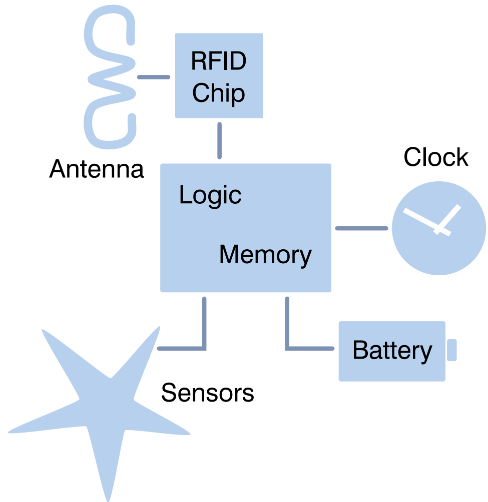

Figure 2: Generic RFID Sensor Package

It is not necessarily required to use battery powered tags and some researchers (notably Philipose et al., [8], and Watters et al., [13]) have used passive tags. Watters et al., focus on structural damage—such as that caused by chloride seeping into concrete, or exceeding temperature limits underneath the heat shield of a space shuttle. In the first case, the sensor makes use of the electro-chemical properties of the medium to produce a voltage change. In the second, the temperature limit is recorded by "blowing a fuse". Both approaches only provide a single bit of information, by preventing the tag from responding. Depending on the application, there may be some difficulty in determining whether the sensor has triggered, or if it was merely that the tag was misread.

Philipose et al., on the other hand, make use of two RFID chips operating on different frequencies so that it is possible to receive a positive signal after the sensor has triggered. Their α-WISP (wireless identification and sensing platform) design uses single bit accelerometers to detect movement. The sensor state is used to shut off either one or the other RFID transponder and thus transmit the sensor state to an RFID reader. Philipose et al., suggest that the transmission of more bits may be possible in the future.

In the middle ground between passive and active lies the *semi-passive* RFID tag. This uses a battery to power the on-board electronics required to run sensors and memory, but uses the passive RFID approach to communicating with the reader.

Sensors may merely provide current state information when the tag is read. Alternatively, sensor data may be logged either on a periodic basis, or variably, depending on when significant events occur.

It may be necessary to equip the tag with a clock so that it can record timestamps. For example, recording the time when a violent shock or other type of event occurs may be important in subsequent diagnosis of the source of the problem.

3.1.2 Logically Connected Sensors

Sensors that are not physically connected may be logically connected. For example, a sensor that measures the temperature of a cold storage area may not be physically wired to tags on items that are placed in the storage area. A logical connection can be set up if we know that an item is within a certain storage area and we know that the temperature sensor is also in that storage area.

Logically connected sensors may still be physically connected to the product and have their own connection to the network, or alternatively they may just be associated with the environment.

This logical connection is partly about knowing the location of the item (e.g. within freezer x) but also about knowing the location of the sensor. In some cases it will merely be necessary to know in which vicinity the sensor resides (e.g. also within freezer x). In other cases, more detailed location information is required (e.g. a laser range finder is at coordinate l, height h and is facing in direction ϕ).

Implicit here is that we need to be able to identify individual sensors. For some sensors, an Internet address will provide sufficient resolution to identify the sensor uniquely. This will be practical where the individual sensor is capable of communicating over the Internet. However, it is not always practical to have the hardware and software to perform this communication for each and every sensor, and in this case, some form of *concentrator* is required. The concentrator would perform the job of supporting the Internet connection of one or more sensors.

In the case where all communication with sensors is over a local network, there may be no need for any form of security. For communicating with external sensors, and for allowing external parties to communicate with internal sensors, it will be necessary to make use of authentication, authorisation, and encryption techniques to ensure that access to sensors is restricted. One way that this might be done is to support a Web Services interface within the concentrator and to make use of the standard security and encryption protocols that go with this [6].

3.1.3 Issues in Connecting Tags to Sensors

There are a number of key issues that we summarise as follows:

- How is battery life affected by physically attaching sensors?

- How much memory is required to store the data associated with a sensor?
- How frequently should data be collected? Should sampling be periodic, event-based, or just record peak values?
- When do we transmit data to the back-end system? Should this data be "pushed" (as it happens) or "pulled" (on demand or during a maintenance operation)?
- What intelligence will exist on-board to process the data "on-line"?

3.2 Sensor Fusion

The term *sensor fusion* refers to a group of computer algorithms or methods that interpret the combined data from different sensory inputs to produce resulting data that is in some sense better or more accurate than if we relied on just a single sensory source [14]. We can think about sensor fusion in various levels of sophistication. In its simplest form, we might simply average the results from a number of individual but similar sensors. In a slightly more complex scenario, we may wish to cope with the possibility that some of the sensors might have failed and thus detect this and eliminate them from the average. Furthermore, it may be the case that the sensors have a non-linear response or the sensors may measure different things. In this case, a simple averaging is no longer appropriate and it is useful to think in terms of a *sensor model* that describes how the different sensors respond to the *state* of the environment. This model might be *stochastic* thus implying that each sensor contributes evidence towards a fact being true, but does not of itself guarantee that fact.

Sensor fusion usually applies to merging information of similar modality, such as the averaging of two temperature sensors. However, it may also apply to sensors with quite different modalities, such as a temperature sensor and a voltage sensor. Both sensors imply something about the underlying state of the product, and this may provide a basis for merging the sensor information.

3.2.1 Issues in Sensor Fusion

- Do we take into account what *actions* have been taken? Actions might be treated as a sensory input.
- Do we fuse data over time, or consider each instant independently? This is similar to the question of whether we use model versus sensor based approaches.
- Are we able to predict future states? It may be useful to consider what sort of decision making process occurs, and to justify any forecasting on this basis.

- What causality relations exist? In particular, how do we model these relationships and what do they imply about the algorithm?
- Is it necessary to take a stochastic or probabilistic approach, or will a deterministic approach be more suitable in some cases?

3.3 Decision Making

The point of most state estimation activities is to allow us to make a decision. Hopefully improving the quality of the estimate of the state will cause a corresponding improvement in the quality of the decision. Ultimately, any decision making scheme must be judged in terms of high-level objectives (such as cost effectiveness, quality of life, etc).

As noted in the section on IVHM, the types of decisions made based on RFID sensor data can roughly be categorised into diagnosis, prognosis, automated inspections and anomaly detection. Diagnosis and prognosis are like two sides of a coin; diagnosis being about events that have occurred, whereas prognosis being about events that might occur in the future. Diagnosis tries to address the problem that any real environment can only be partially observed. By looking at the history of events it is sometimes possible to identify a root cause. For example, if a radio oscillator has been found to have drifted, a record of shocks that have occurred to it might help identify if this was caused by a heavy landing or some other factor.

Prognosis is about predicting the future state of an object. In particular, it is about using the prediction to make a better decision now. An example is the shift from performing preventative maintenance, which tends to assume that all objects of a type wear at the same rate, to performing *predictive* maintenance, which looks at actual usage to estimate when the part will wear out, and to schedule maintenance to occur prior to this.

Automated inspection can be seen to sit in between prognosis and diagnosis, referring to the *current* state rather than the past or future state. A typical example is to use a special sensor on welded joins. Welds tend to be the weakest part of a structure and will be the first to break. A sensor that changes its electrical properties as the weld deteriorates can be used to obviate the need for other, more expensive tests. By incorporating such a sensor in an RFID package, regular maintenance checks can be performed easily and cheaply. An additional benefit of this approach is that it provides an audit trail of when each inspection was performed.

The final decision type is that of anomaly detection. This is similar to diagnosis, in that the analysis is about a past event. In this case, however, the tag is taking an active role in identifying that there is cause for alarm, and sending an alert. A difficulty here lies in the problem of false alarms. The tag should not produce too many unnecessary alert messages or else there will be the tendency for the human operator to ignore or disable the tag.

3.3.1 Issues in Multi-Sensor Decision Making

In terms of this project, there are a number of key issues:

- What sorts of decisions need to be made? For example, we might be attempting to predict when to service. Alternatively, we may want to direct the servicing operation better.
- How can we address the issue of false alarms in anomaly detection?
- Should decision making be kept separate from state estimation or be combined in some way? For example, there may be decisions that we take to find out more information.

4 Proposed Methodology

In the previous section, the area of RFID sensor integration was discussed, both in terms of possible technological issues, and in terms of applications that have been or might be used. In this section, we propose a methodology for studying this area in more detail. The methodology is summarised in the diagram in figure 3. These five stages are now discussed in more detail.

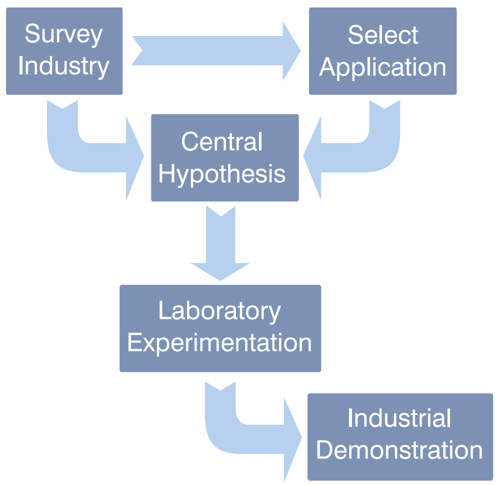

Figure 3: Diagram showing the five process steps in proposed methodology

Published February 1, 2006. Distribution restricted to sponsors until August 1, 2006

4.1 Surveying and interviewing industry partners to help identify key areas

This task will use surveys and interviews to establish more detail about specific applications. It is expected that the output of this section will be in terms of key problem areas that should be the focus of subsequent work.

In addition to surveys and interviews, it may be possible to capture real sensor data for one or more of the proposed application areas. This type of data, if it is possible to capture it, will be useful in testing theories and validating any simulation or laboratory experiments.

4.2 Selecting an application area

A key issue that has shown up early in our communication with industrial partners is that there are many possible application areas that might be discussed, but not all should necessarily involve both RFID and sensors, nor require the integration of the two. During those first discussions, we arrived at a set of questions to try to identify the most relevant applications. These were:

- Is radio frequency communication required? (Or is there always some form of wired connection.)
- Is identity important? (Or is this mainly a sensor application.)
- Have other forms of wireless connection been considered? (Such as Bluetooth, or Zigbee.)
- Is the sensor on the physical object in question or is the sensor in the environment?

Following this, we realised that mere yes / no answers provided little disambiguation of the different application sectors. For this reason, we are in the process of developing a set of characteristic axes. The axes measure features such as the frequency of communication or the complexity of the sensory information. Different applications will be plotted on these axes and we will then attempt to identify popular clusters or common themes. The results of this work will guide subsequent research.

4.3 Proposing one or more approaches to address the identified problems

This task leads on from the previous two, using the identified key problem areas and selected application group, and coming up with a proposed approach. This approach will be in terms of algorithms, tools, or methods for solving the identified problems. At this stage, a detailed description of the approach will be generated along with a draft of our central hypotheses.

4.4 Testing the approaches through simulated and / or laboratory experiment

This section aims to test the hypotheses generated in the previous section. It is expected that this will involve either some form of simulation or physical experiment. The results from this section will be used to show whether or not there is a clear benefit in using the approach.

4.5 Dissemination of results to industry partners

Although this final stage may seem obvious, it is important to emphasise that the useful results of this work must be disseminated widely amongst industry partners in the Aero-ID program. Furthermore, this stage should aim to convert the laboratory and theoretical work into an industrial demonstration. It is expected that the industrial partner will take a lead role in this process.

5 Conclusions

The overall aim of the RFID Sensor Integration theme is to investigate the idea of integrating RFID and other types of sensory data in the context of the aerospace sector. This specific area has not been widely studied, but similar areas involving sensor fusion suggest that it is likely to be a rich source of valuable insights.

This overall objective, however, is quite broad and it is necessary to focus the work in a number of ways. First and foremost, it should be focused to be applicable to industry and to areas of popular industrial concern. The first two stages of the methodology should help to ensure this. Second, it should fit well with the skills and expertise of the researchers. Finally, while being ambitious, it should be sufficiently tightly defined so that it can be completed in the time available.

Acknowledgments

The authors wish to thank Arun Ayyagari for valuable insights, and also Lila Theodorou, Duncan McFarlane, and Victor Prodonoff Jr., for useful comments on draft versions of this report.

References

[1] Anon. Clinisense home page, 2006. Available from http://www.clinisense.com [Online; accessed 27-January-2006].

[2] Anon. Infratab home page, 2006. Available from http://www.infratab.com [Online; accessed 27-January-2006].

[3] Jonathan Collins. Sensing a product's shelf life. *RFID Journal*, April 2005. Available from: http://www.rfidjournal.com/article/articleview/1539/1/1/ [Accessed 27-Jan-2006].

[4] Dimitris Kiritsis. Ubiquitous product lifecycle management using product embedded information devices. In *Proc. IMS 2004*, 2004. Available from http://www.promise.no [Online; Accessed 29-January-2006].

[5] Robin R. Murphy. Dempster-shafer theory for sensor fusion in autonomous mobile robots. *IEEE Transactions on Robotics and Automation*, 14(2):197–206, 1998.

[6] Anthony Nadalin, Chris Kaler, Philip Hallam-Baker, and Ronald Monzillo. Web services security: SOAP message security 1.0 (WS-security 2004). Technical report, OASIS, 2004. Available from: http://docs.oasis-open.org/wss/2004/01/oasis-200401-wss-soap-message-security-1.0 [Accessed 19-Jan-2006].

[7] Stan Ofsthun. Integrated vehicle health management for aerospace platforms. *IEEE Instrumentation & Measurement Magazine*, 5(3):21–24, Sept. 2002.

[8] Matthai Philipose, Joshua R. Smith, Bing Jiang, Alexander Mamishev, Sumit Roy, and Kishore Sundara-Rajan. Battery-free wireless identification and sensing. *IEEE Pervasive Computing*, 4(1):37–45, January–March 2005.

[9] Stuart Russell and Peter Norvig. *Artificial Intelligence: A Modern Approach*. Prentice Hall, second edition, 2003.

[10] Alessandro Saffiotti. The uses of fuzzy logic in autonomous robot navigation. *Soft Computing*, 1(4):180–197, December 1997.

[11] Vince Stanford. Pervasive computing goes the last hundred feet with RFID systems. *IEEE Pervasive Computing*, 2(2):9–14, 2003.

[12] Ken Traub, Greg Allgair, Henri Barthel, Leo Burnstein, John Garrett, Bernie Hogan, Bryan Rodgrigues, Sanjay Sarma, Johannes Schmidt, Chuck Schramek, Roger Stewart, and KK Suen. The EPCglobal architecture framework, July 2005. Available from http://www.epcglobalinc.com/standards_technology/Final-epcglobal-arch-20050701.pdf [Online; accessed 29-January-2006].

[13] D.G. Watters, P. Jayaweera, A.J. Bahr, and D.L. Huestis. Design and performance of wireless sensors for structural health monitoring. In *AIP Conference Proceedings*, volume 615, pages 969–976, May 25 2002.

[14] Wikipedia. Sensor fusion — wikipedia, the free encyclopedia, 2006. Available from `http://en.wikipedia.org/w/index.php?title=Sensor_fusion&oldid=34143496` [Online; accessed 27-January-2006].

A Algorithms for Sensor Fusion

This appendix provides some further background on algorithms used in sensor fusion, which may be of interest to the more technically inclined reader.

Some of the typical approaches used to perform sensor fusion include Kalman filters, Bayesian Nets and Dempster-Shafer theory. A Kalman filter provides a way of estimating the state from a time series. Roughly speaking a Kalman filter is based on assuming that the next state is a linear function of the current state plus some Gaussian noise. This assumption is usually quite appropriate for many tasks but does mean that only a single hypothesis can be tracked. Multiple hypothesis techniques such as Markov Chain Monte Carlo (MCMC) may be more appropriate for some problems. Note that Kalman filters have the advantage that they are able to deal with continuous variables. Many other techniques can only work with continuous variables after some sort of discretisation.

Bayesian Nets are a general class of algorithms based on the idea of reasoning about the conditional probabilities between state variables. The simplest form is a Markov Chain (or first order Markov Process), which is a simple time series of variables with any variable at time t being conditionally dependent on (and only on) the variable at time $t-1$. In general, Bayesian nets can include much more complex sets of conditional interdependencies. Bayesian nets are generally based on a user specified model of where the conditional dependencies lie (and possibly what the conditional probabilities are—although this can be derived from data in some cases).

The use of Dempster-Shafer theory for sensor fusion has been explored by Murphy [5]. Dempster-Shafer uses *belief* as the measure of likelihood of a fact, rather than using a probability measure. The key problem for the use of DS theory and other forms of possibilistic logics, as Murphy points out, is that sensor occlusion or failure is not dealt with properly. Murphy shows, however, that DS theory can be adjusted to allow for such things as sensor failure by explicitly representing ignorance.

Fuzzy Logic has long been used in embedded environments as a way of representing the uncertainty in sensors without the complexity (and perhaps intractability) of probabilistic measures. Some of the issues around using Fuzzy Logic (particularly in a control context) are given by Saffiotti [10].

There is a strong argument for discarding logics based on measures other than probability. Russell and Norvig [9, page 474] summarise Bruno de Finetti's argument that concludes that it is irrational to accept beliefs that violate the axioms of probability. Bayesian methods may be preferred from this point of view, but may also be too complex to calculate in an embedded processor. Increases in the computational capabilities of on-board computers, as well as improvements in the efficiency of Bayesian algorithms seem to be leading to a corresponding rise in popularity of these methods.