Published November 1, 2002. Distribution restricted to Sponsors until February 1, 2003.

AUTO-ID CENTRE

WHITE PAPER

The Need for a Universal Smart Sensor Network

Stephan Haller, Steve Hodges

AUTO-ID CENTRE INSTITUTE FOR MANUFACTURING, UNIVERSITY OF CAMBRIDGE, MILL LANE, CAMBRIDGE, CB2 1RX, UNITED KINGDOM

ABSTRACT

The benefits that tagging technologies can bring to enterprises by bridging the gap between the real world and its representation in information systems has clearly been demonstrated by the work of the Auto-ID Centre and in many case studies performed by different companies.

While RFID tagging currently receives a lot of attention, it is only one of many different technologies that may be used to get information about objects in the real world. In this paper, we address this much larger set of 'smart sensor' technologies. Each of these technologies typically has its own particular strengths and weaknesses in different application areas and under different operating conditions. Therefore, for many applications it seems sensible to integrate the data from a combination of underlying sensors to form a smart sensor network.

This paper describes the underlying technologies that may be used to build a smart sensor network and presents the authors' vision of the network itself, providing several example scenarios to demonstrate its utility. Potential problem areas are also discussed. A framework for a universal smart sensor network is proposed; without such a framework, the effort of integrating all the different sources would be prohibitive, resulting in the continued proliferation of inflexible systems that are unable to adapt to the fast changing business requirements of today's world.

We believe that the time is right for further study of the technology behind smart sensor networks and the applications for them.

Published November 1, 2002. Distribution restricted to Sponsors until February 1, 2003.

WHITE PAPER

The Need for a Universal Smart Sensor Network

Biographies

Stephan Haller
Senior Scientist, SAP AG

Stephan Haller leads the Smart Items research program efforts at SAP Corporate Research in Karlsruhe, Germany. He has been working with RFID Technology since 1999 in his previous capacity as a Software Engineer and Development Manager at SAP Labs in Tokyo. Before joining SAP Japan in 1997, Stephan worked for 3 years as a research engineer at Matsushita Electric Works' Central Research Laboratory in Osaka, Japan, developing distributed information systems for building control and home health care. He holds a Master's degree in Computer Science from the Federal Institute of Technology (ETH) in Zürich, Switzerland.

Steve Hodges
Associate Director, Auto-ID Centre Europe

Steve Hodges is an Associate Director at the Auto-ID Centre Lab at Cambridge University. Steve received his first degree in Computer Science with Electronic Engineering, from University College London, and received his Ph.D. from Cambridge University Engineering Department in the area of Robotics and Computer Vision. His interests include embedded sensor systems, intelligent devices, computer augmented environments, mobile robotics, image processing, low-power radio communication, mass customization of consumer products and RF tagging technologies.

WHITE PAPER

The Need for a Universal Smart Sensor Network

Contents

1. Introduction 3
   1.1. What is Auto-ID? 3
   1.2. What are Smart Sensors and Smart Sensor Networks? 3
   1.3. Background of this White Paper 3
2. Smart Sensor Technologies 4
   2.1. Identification Technologies 4
   2.2. Location Sensing 5
   2.3. Sensing the Physical Environment 5
   2.4. Keeping Track of Time 6
   2.5. Communication 6
3. Smart Sensor Applications 8
   3.1. Applications in Use Today 8
   3.2. Envisaged Applications 8
4. Vision of a Universal Smart Sensor Network 11
   4.1. A Distributed, Agent-Based Approach 11
   4.2. The Abilities of a Smart Sensor Network 12
   4.3. The Time is Right 12
   4.4. Smart Sensors in an Auto-ID Framework 13
5. Problem Areas 13
   5.1. Standards 14
   5.2. Collaboration 15
   5.3. Software Frameworks 15
   5.4. Hardware 16
6. Conclusion 17
7. References 18

1. INTRODUCTION

1.1. What is Auto-ID?

The term 'Auto-ID' is ambiguous; it is often used by different people to mean different things in different contexts. The simplest interpretation of an 'Auto-ID system' is one that may be used to automate the identification of objects. A number of underlying technologies could be used to implement such automatic identification, and the characteristics of these – such as accuracy, performance and cost – vary quite considerably. However, the term 'Auto-ID' is sometimes also used to refer to much more sophisticated systems deployed to (automatically) glean other types of information about objects, beyond a simple object identifier.

Finally, 'Auto-ID' is also used as shorthand for 'The Auto-ID Centre', such that references to 'Auto-ID systems', 'Auto-ID tags', 'Auto-ID readers' and so on refer very specifically to work associated with the Auto-ID Centre in these areas. This is exactly the meaning ascribed to 'Auto-ID' throughout this paper.

1.2. What are Smart Sensors and Smart Sensor Networks?

A smart sensor is, most simply stated, a device that
1. knows some information about itself (or the object with which it is associated) and
2. has the ability to communicate this.

By this definition a radio frequency identification (RFID) tag that contains information about the object it is attached to is, at the simplest extreme, a smart sensor. But smart sensors can also be much more sophisticated than RFID tags in their design and functionality – they may be able to sense things in their environment; they may contain very detailed information about an associated object; and they may be able to communicate with each other autonomously.

Indeed, the ability of smart sensors to communicate with their peers leads to the notion of smart sensor networks, where information can be passed between the network nodes (the smart sensors themselves) and decisions perhaps made automatically. A very simple network could be formed when a number of smart sensors are in the presence of a co-ordinating reader, which can effectively act as the infomediary that gleans the relevant information from the sensors within range and makes the decisions on their behalf. In this way, even basic RFID tags can form a smart sensor network of sorts. At the other end of the spectrum, more sophisticated smart sensor nodes would carry their own power source, and would be capable of exchanging information with their peers and making decisions internally.

1.3. Background of this White Paper

The Auto-ID Centre is promoting a number of technologies centred around the automatic identification of objects that may be combined to provide previously unheard of visibility of objects in the supply chain and beyond. The Centre believes that this new ability will deliver tangible business benefits to a wide range of organisations. Within the Centre, a number of studies that demonstrate the specific benefits of the use of these technologies have been carried out [kambil2002, alexander2002a, alexander2002b], and more such studies are underway.

Additionally, many studies carried out independently in recent years also indicate the business benefits that may be achieved through the deployment of systems that can automatically identify assets as they move along the supply chain.

To date, most of these studies have limited their consideration to the simplest form of smart sensors – namely simple identifiers attached to objects. We believe that there are many more benefits to be gained through the use of more sophisticated sensor networks. In our vision, supply chain visibility is much more than simply knowing where objects are – many more characteristics of those objects will be known, even as they change over time.

This paper provides an introduction to the features we might expect of smart sensors and the networks they form, both using today's technology and also in the future. It presents some of the problems that must be overcome to make this sensing suitable for deployment, and it looks at some possible usage scenarios and the benefits that may be gained. We believe that the time is right for further study of the technology behind smart sensor networks and the applications for them.

2. SMART SENSOR TECHNOLOGIES

Having defined a smart sensor as a device that knows something about the object that it is associated with, and is able to communicate this, this section reviews the various underlying technologies that may be used as a foundation for building smart sensors.

2.1. Identification Technologies

As mentioned in Section 1.2, one of the simplest forms of smart sensor, by our definition, is a device that contains some fixed information about the object that it is associated with. Examples of such information include the stock keeping number of a product, or the shipment serial number of a package. In these simple examples, the device may take the form of a tag of some kind that is physically attached to the associated object. Two common tagging technologies are barcodes and RFID tags, which are discussed in more detail below, along with a third technique that we have classed as 'fingerprinting'.

2.1.1. Barcodes

The familiar black-and-white stripes found today on most consumer packaged goods contain information describing the type of an item and the identity of its manufacturer. This information is communicated through the use of a barcode scanner, an optical device that can read such a barcode when it is held in the reader's field of view. In addition to the commonplace one dimensional (1D) barcodes, a number of more sophisticated symbologies capable of representing additional information have been developed. These include the 1½D or stacked barcode, and the 2D barcode [aim, barnes1999].

The barcode standards most commonly used for consumer item labelling dictate that the identity of the manufacturer and the stock-keeping unit (SKU) of each item is identified in two separate sections of the barcode. Although barcodes can theoretically be of any length, in practice they are limited in size and contain no serial number.

2.1.2. Radio frequency identification tags

Frequently deployed in vertical application areas, RFID tags typically combine a modest storage capacity with a means of communicating this stored information wirelessly to an RFID reader. Crucially, unlike barcodes, line-of-sight is not needed for this information transfer. A large number of different RFID technologies are in development or deployed in specific niche applications, but the most common approach to implementing an RFID tag is to mount a tiny silicon chip onto a foil antenna – the resulting tag often takes the form of a label. Such RFID tags are passive (i.e. require no on-board battery) and can typically be read from a distance ranging from a few centimetres to a few metres. There are also

active tags which do have an on-board battery, resulting in far longer read ranges and memory sizes, but at far higher cost and a limited lifetime of 2–5 years. For more details on the vast range of tag characteristics possible, the reader is referred to [finkenzeller1999].

Typically RFID tags contain more storage capacity than barcodes, and it is therefore possible to further partition the memory to include a serial number, so that items can be uniquely identified. It is also possible to design an RFID tag with read/write capabilities, such that some or all of the information it contains can be modified as required.

2.1.3. Fingerprinting techniques

Whilst not strictly a form of smart sensor tagging, there are a number of fingerprinting techniques in development which are mentioned here for the sake of completeness. Rather than associating some particular information or data with an object, it is possible to use some of the physical properties of that object to generate a unique 'fingerprint' of the object that may be used to uniquely identify it [pappu2002]. This unique identity may then be used as a key to refer to more specific information elsewhere, much like the electronic product code (EPC™) [brock2001a] is used in an Auto-ID system.

2.2. Location Sensing

In addition to fixed information about an object, such as its identity, there is a lot of dynamic information associated with objects – information that changes in real-time. A good example of this is the current location of an object, which will obviously change whenever the object is moved. Depending on what this location information is used for, the nature and granularity required may vary dramatically. The nature of location information may be split into a number of broad categories. In addition to **absolute** location information, where a co-ordinate system of some kind is used, it is also possible to describe location in terms of **containment** (when for example an object is known to be in a certain warehouse) or in terms of **proximity** or **range** (i.e. distance) to either a fixed point or to other mobile objects. The nature of information available will tend to depend on the underlying technology used. For example, the global positioning system (GPS) [leick1995] generates absolute location data (in the form of GPS coordinates), whereas RFID tag readers generate proximity information.

In addition to the fundamental nature of the location information, the granularity (or resolution) as well as the accuracy and the speed of update of the information must also be considered. Different technologies also have different requirements in terms of the supporting infrastructure needed, the implementation cost, power consumption, component size and so on.

2.3. Sensing the Physical Environment

A smart sensor may measure different characteristics of the physical environment that it (and the object with which it is associated) is exposed to. Examples include:

- temperature;
- pressure;
- humidity;
- shock/vibration/acceleration.

2.4. Keeping Track of Time

Any of the above sensing technologies will generate data that varies with time, because the sensed characteristic will itself tend to vary with time. If a sensor has the ability to keep track of time, and also has some storage capacity, then it is possible to keep track of these changes. The resolution and frequency of readings stored would depend on the amount of memory available, and also on the period for which this information needs to be stored on the device (before being downloaded to the network). It may not be necessary to log sensory readings at a fixed interval; additional intelligence built into the sensor might dictate that only significant changes in the environment be stored in memory in this way.

To differentiate a **simple** smart sensor from one that has the ability to periodically store readings and timestamp information, we refer to the latter as a **logging** smart sensor.

2.5. Communication

The definition of a smart sensor in Section 1.2 referred to the ability of the sensor to communicate information. This section briefly reviews some of the relevant communication technologies.

2.5.1. Light

The simple barcode communicates its fixed information payload by spatially modulating incident light – parts of the code reflect and other parts do not. A barcode can therefore be read simply by shining light at it, either uniformly or as a scanned beam, and looking at the intensity of the reflection. Most barcodes work well with either visible or infrared light.

Another common form of digital communication using light is the infrared data port (or IrDA port) present on many electronic devices, such as mobile telephones, electronic organisers and laptop computers. High data rates over relatively short distances are possible using IrDA communication [irda], and the IrDA components themselves are relatively low-powered during operation, making them a candidate for smart sensor communication.

The downside to light communication systems is the requirement for line-of-sight between the communicating devices. That is, the transmitter and receiver components must have a direct, unobstructed straight-line path between them for the duration of the communication.

2.5.2. Active, short-range radio communication

Unlike communication using light, radio communication works without line-of-sight – in some cases there can be quite substantial obstacles between the transmitter and receiver and communication is still viable. (The details will vary greatly depending on exactly how the communication is set up.) This clearly is a great advantage when the physical environment is less well constrained, and line-of-sight cannot be guaranteed.

As the term implies, active radio communication requires a power source on each side of the link, and typically this is quite significant – a small battery will not be able to power such a system for long periods (e.g. many months) without replacement or recharging. Short range radio devices (also called low power radio devices) can communicate over distances of up to around 100 metres, and are often designed to support direct communication between peers.

In addition to proprietary solutions, there are a number of standards that might be suitable for active, short-range communication between smart sensors. These include Bluetooth [miller2001] and IEEE 802.15.4 (a.k.a. Zigbee) [gorday2001].

2.5.3. Active, cellular radio communication

When radio communication over distances longer than around 100m is needed, there are basically two approaches that may be adopted. The obvious approach of increasing the RF transmission power has a number of disadvantages (not least of which is the electrical power required). The alternative approach, which has enjoyed particular popularity in deployment of cellular telephony and WiFi wireless networking, is to build an infrastructure of radio basestations such that communication only needs to occur to the nearest basestation. The required communication between basestations is implemented over a standard wired network. Once again, there are many characteristics of such a cellular set-up, but a number of standards exist, including GSM, GPRS, UMTS, WLAN (a.k.a. WiFi) [oliphant2000, riezenman2002].

2.5.4. Passive radio communication

In addition to active radio communication, where both ends of the communication channel are powered, it is possible to design a passive communication system, where only one of the endpoints (the master) has an external power source, which is used to power the entire communication. This is exactly how passive RFID tags are read – the reader is a master that generates an RF field that powers the tags (the slaves) within range in addition to transmitting data to those tags.

Even when a device does have its own power source, it may be advantageous to use such a system to provide a zero-power communications channel. In this way, none of the device's precious batteries need be used when communicating – and this can be very important due to the relatively high energy demands of radio transmissions and reception (see Section 2.5.2). This is exactly the approach taken in the design of semi-passive RFID tags – the communication occurs with no power being supplied by the tag, but an on-board battery does enable other tag functions.

It is also possible to imagine a scenario where a particular device operates as a master for some of the time, and as a slave at other times, depending on circumstances.

An initiative to define a standard for such passive radio communication, Near Field Communication, has recently been announced [philips2002].

2.5.5. Wireline communication

Whilst smart sensors will often be mobile, wireless devices that rely on some kind of wireless communication, it is also possible to imagine fixed sensors. For example, a temperature sensor in a warehouse may be permanently deployed in a fixed location so that it can provide information about the ambient temperature in that location and so that the system may infer how that reflects in the temperature of items stored there.

In many cases, it will be convenient to hard wire power and communication services to such fixed sensors, resulting in higher performance, robustness, longevity and so on. There are many mature wireline communication systems and protocols, but of particular relevance to the requirements and philosophy of the smart sensor network are recent developments of the concept of **network-direct** appliances. This approach involves incorporating a basic ethernet interface into electronic devices, even when they are small, cheap and embedded in operation. (In the past, the ethernet interface has been restricted to more sophisticated electronic systems, such as personal computers and

workstations.) A lightweight TCP/IP stack is typically implemented within the network-direct appliance, so that it can communicate over the ethernet with pretty much any other device on the Internet. This clearly builds incredible flexibility into the system, and for this reason it may well be a suitable technology for use in a smart sensor network.

3. SMART SENSOR APPLICATIONS

3.1. Applications in Use Today

The various forms of smart sensor technology have been deployed in different vertical application areas for many years; for example RFID tagging technology is commonly used as part of a proprietary, bespoke solution for a given industrial problem. Although enterprises often use up-to-date technologies, each implement their own proprietary solution; they may even have multiple incompatible solutions within the same enterprise.

Such an approach results in several problems:

- the solution is expensive to maintain and even harder to adapt to new business process requirements;
- it doesn't work across enterprises;
- the advantage may be optimal for given applications, but as technology advances the need for optimality is removed, and the advantages of interoperability become compelling.

Just as the Auto-ID Centre is building common, open standards and infrastructure for the use of RFID tags, there is a need for common, open standards and infrastructure for sensor networks in general. The benefits of smart sensors will be much greater if they interoperate seamlessly, and if smart sensor networks can be built from the sensor components with very little effort. Ideally this would actually be an automatic process.

3.2. Envisaged Applications

To date the full potential of smart sensor networks is far from realised. We imagine a great many diverse sensors embedded throughout certain environments, continually monitoring and communicating to enhance the abilities of our systems and to save wasted effort and money. An early version of such an application is described in [thede2001]. To give a flavour of the applications that may be commonplace in the future, we present a number of example scenarios where smart sensor networks play a pivotal role:

- The temperature of frozen and chilled goods can be monitored continually throughout the cold chain, to ensure that the relevant standards are adhered to. In the longer term, it may be possible to introduce more sophisticated algorithms for calculating sell-by and use-by dates, based on the specific conditions experienced by individual products. Taking this idea through to its logical conclusion, dynamically updated sell-by and use-by dates may become commonplace, not just within the cold chain, but even through to the consumer, enabled by smart sensing in the home. A similar scenario from the pharmaceutical area is depicted in Figure 1.

Figure 1: Example of a smart sensor network based pharmaceutical supply chain.

The products are packaged in smart blister packs that can sense the current temperature in addition to identifying the product. Environmental temperature is controlled at all points in the supply chain. For example, smart trucks contain refrigeration units that are linked into the smart sensor network for control of temperature during transportation. Similarly, there are smart racks, shelves and cabinets. Data about the blister packs is sent to the manufacturer for tracking and replenishment calculation purposes. At each location, the expiry dates can be recalculated. There are two ways of doing this: either all temperature data is sent back to the manufacturer, where the updated expiration dates are calculated and stored for subsequent access from remote locations. Alternatively, the manufacturer provides a calculation algorithm in the associated physical markup language (PML), which can be downloaded and executed at the blister pack's remote location.

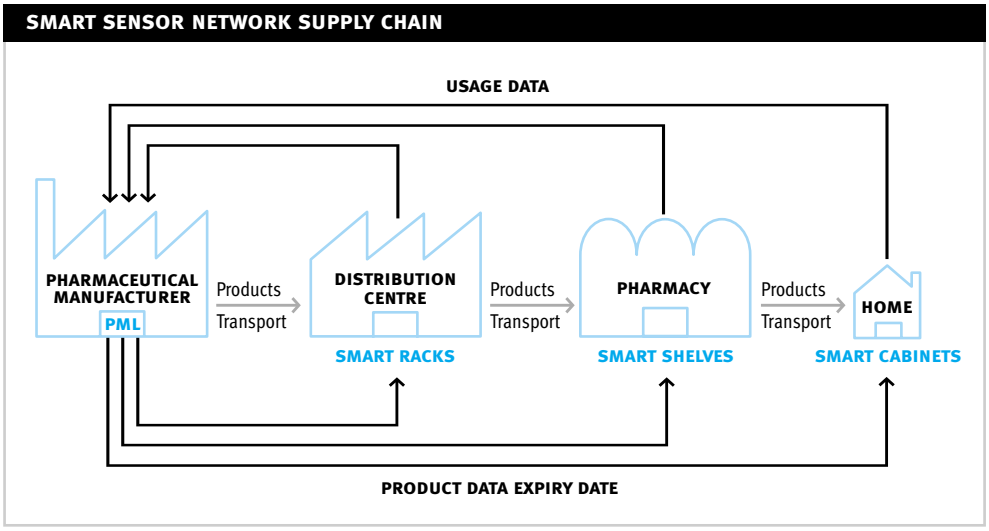

– Goods sensitive to shock and vibration (for example electronics goods and sensitive equipment) can be monitored throughout transit and storage to ensure that their quality is not compromised. Any potentially harmful treatment will be detected and subsequently flagged, making sure that defective goods are taken from the supply chain as early as possible and thereby saving the costs associated with failing to do this. Additionally, the source of the shock/vibration will be easier to track down (for example the time and location of the problem) using the data generated by the smart sensors, and this will facilitate remedial action to prevent reoccurrences.

– In addition to temperature monitoring, other environmental factors such as pressure and humidity can be detected and monitored by smart sensors. This may have implications for the safe transit and storage of many different items, including foodstuffs, electronic components and chemicals.

Figure 2: Smart Truck

1. Wireless Communication
2. GPS
3. Temperature RFID on blister packs

Inside the truck are one or more readers that can interrogate the temperature-sensing RFID tags continuously. The information is collected and together with GPS location information sent to backend systems via some form of wireless connection. Note that during transit, the truck will not be online all the time, both for connectivity as well as cost reasons, but data collection continues. At certain points (sometimes also referred to as "info stations" [badrinath1996, ye1998, kubach2001]) during transit connectivity is re-established and synchronisation is done automatically. Note that the truck is a small smart sensor network itself, and while connected it is a node in the universal smart sensor network.

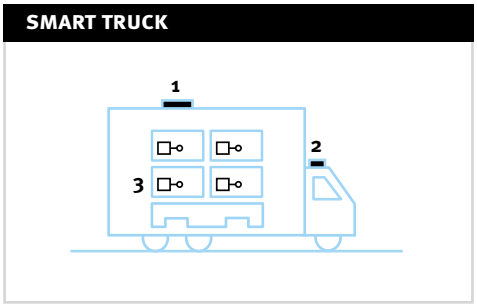

– Vending machines will automatically know the sell-by dates of the products they contain, and will therefore be able to reduce the prices of relevant goods to reduce wastage; these decisions will either be made independently, or may involve collaboration with other machines or with manufacturers and retailers. The machines may also be aware of other factors, such as the re-stocking schedule and environmental conditions (e.g. the temperature, which has a big effect on the demand for chilled drinks), and may use this information to influence decision-making. Of course, these ideas are not just limited to drinks vending machines – in parts of the world where a whole host of vending machines are commonplace, such as Japan, the potential applications of smart sensors are even greater.

– Recent developments in olfactory sensor devices open up the possibility of freshness monitoring of foodstuffs based on smell. It may be appropriate to combine the data generated by a number of heterogeneous sensors in order to provide a more reliable estimate of freshness, but this task would be ideally suited to a smart sensor network.

– In addition to associating basic identifiers with objects through tagging to enable tracking throughout the supply chain, it is also possible to store additional data on such smart sensor tags. This may make operations in a network-free environment considerably easier. For example, when performing maintenance at remote sites such as oil rigs, electricity distribution masts or even countryside garages [chaxel1997], it may not be possible to retrieve and update records that are stored on a networked enterprise database – instead it may be much more appropriate to store the information with the objects themselves. Even when network access is technically possible, it may be appropriate to reduce the reader and networking infrastructure required to reduce cost, ease deployment, or reduce power requirements.

– Smart sensors can be used to build event and error detection and logging capabilities into devices to make them more intelligent in the face of unforeseen circumstances. It may be possible to run diagnostic operations to detect certain operating characteristics or error conditions, and either log this information for subsequent reference, or communicate it in real-time across a smart sensor network. In this way, it may be possible to build up a much more detailed understanding of the typical usage and behaviour of machines and systems, and even to anticipate fault conditions.

– In applications where secure information transfer is required, smart sensor devices with security built in will ensure that privacy and authenticity is guaranteed [sarma2002]. In fact, these features may quickly become commonplace in smart sensor networks to ensure the integrity of enterprise applications built on top of the technologies discussed in this paper.

– Smart sensors in machines and operating equipment will record the exact usage and operating parameters like temperature, time of idleness, peak loads etc. Feeding this information back to product design will lead to improvements in the next version of the machine. The product can be designed to match the actual utilisation more closely, resulting in lower operating costs because of lower maintenance costs or more efficient power usage. Also the handling of the machine can be simplified if rarely used functions are eliminated and often-used functions are made easier to use.

By providing a single platform that encompasses all the technologies discussed in Section 2, a universal smart sensor network would facilitate the deployment of all the scenarios provided above, as well as others. It would be much easier to deploy systems that exhibit sophisticated behaviours.

This platform would avoid the ‘media breaks’ [fleischo1] introduced when data is transferred between different media – e.g. reading the shipment number from a label and manually typing the information into a computer system. It also alleviates the problem of integrating different proprietary systems in a single application, thereby avoiding the interfacing effort that is usually required and also guaranteeing the correctness, consistency and completeness of the data. As a result, applications ranging from environmental sensing and information gathering through to those exhibiting automated control of the environment may be built much more readily.

4. VISION OF A UNIVERSAL SMART SENSOR NETWORK

Our vision is for an environment where embedded smart sensors of all types are pervasive. Although these devices would range greatly in their complexity and ability, they would communicate with each other and with infrastructural components (such as network interfaces and processors) seamlessly whenever appropriate. Should parts of the sensor network become disconnected, then operation would automatically continue within the resulting partial networks for the duration of the disconnection, and information would be resolved seamlessly when full connectivity is resumed. New devices and new technologies could be integrated into the network autonomously, by means of a description and discovery protocol that 'matches up' devices that can usefully communicate with each other, providing true "plug'n'play".

Information would by its very nature be timely and up-to-date, and would be passed around this smart sensor network automatically, as appropriate. Intelligent software agents would be able to interface to this network and would make decisions based on the information, thereby automating control of the environment.

Being able to track every object throughout its entire life-cycle, potentially including the period when a consumer is using it, has severe privacy implications. The authors feel that these issues are, however, outside the scope of this paper. They have been addressed elsewhere in the literature [langheinrich2002], and will undoubtedly be given much more consideration in the future.

The rest of this section discusses how our smart sensor network vision might be realised in terms of an implementation approach, and what the abilities inherent in such a smart sensor network might be. We explain why we think that now is the right time to be exploring this vision and why we think that the Auto-ID Centre framework is particularly relevant.

4.1. A Distributed, Agent-Based Approach

The power of modularity in software systems is well known [booch1994]. Software modules, also referred to as software components, allow sophisticated and flexible systems to be constructed with much reduced effort when compared to the construction of an equivalent bespoke, monolithic solution. Each of these components provides a certain functionality or set of services, has a well-defined interface and performs well-defined operations. Typically, the components run together in a special environment on a computer system, the operating system, which provides core low-level services such as memory management and inter-component communication. Something similar to a component-based methodology may be appropriate for building a smart sensor network. Each sensor device is represented by its own software component, or more accurately, a software agent. These agents act on behalf of the associated devices. Each has a well-defined interface and different agents can be plumbed together using services provided by an underlying smart sensor network operating environment. This same operating environment can also provide a mechanism for registration and discovery of sensors within the network. The actual run-time code that embodies each software agent could either be embedded into each device, or it could be made available from a server on the smart sensor network. Similarly, it may actually be executed within the sensor device, or it may be run on the network through a service provided by the operating environment.

One of the advantages of an agent-based approach is that it lends itself to a distributed implementation. An application is made up of a number of agents that interact with each other, and each of these may run on different platforms, such as the sensor devices themselves. This naturally provides all the advantages associated with a distributed approach, such as fault tolerance, location transparency and scalability [coulouris1988]. In fact, the approach of representing different devices using agents and allowing them to interact across a network is well known by distributed systems programmers as agent-based design

[luck2001, wooldridge1997]. When the software agents are associated with corresponding specific devices, this is sometimes called a holonic system [bussmann1998, mcfarlane2000].

4.2. The Abilities of a Smart Sensor Network

The smart sensor networks that we envision will support a number of novel and interesting features; some of these are listed below:

- Multiple homogeneous smart sensor devices may readily be combined to overcome the limitations of single sensors. The benefits of this are multiple – the redundancy of information allows more accurate sensing of the environment and also raises the possibility that faulty or mis-calibrated sensors can be detected. Also, spatial distribution of sensors allows a much more complete view of the environmental conditions to be built up.

- Integration of heterogeneous sensing technologies will also be possible. For example, whilst a temperature sensor with an inbuilt real-time clock can be used to monitor perishable foods and estimate sell-by and use-by dates, it's possible that the combination of temperature and smell sensing would provide a much more robust detection of end-of-life.

- Sensors designed for different applications areas and/or built by different technology vendors will be interoperable, making the deployment of real systems much easier. Indeed, we imagine flexible solutions that can be re-configured (even on-the-fly) to optimise processes with respect to changing variables and requirements.

- Smart sensor networks will scale dynamically and automatically. For example, the scope of a particular scenario may be limited to a single company; however, in a different application, the sensor network might involve the entire supply chain.

- As business requirements change, it is relatively straightforward to modify the underlying processes since the sensing and communication aspects of the systems are very flexible. When a new configuration is created, the smart sensor devices will detect this and automatically operate in a useful way.

- As new types of sensor devices are developed, they can readily be integrated into the smart sensor network, enabling the benefits to be captured quickly and facilitating experimentation. Similarly, new combinations of sensors and various locations and configurations of sensors can be prototyped much more easily.

- Data regarding the actual usage of products and equipment may automatically be collected continuously. In addition to providing an insight into the processes within an organisation, this information could also be used to help improve actual product design, by learning from how products are actually used (rather than how they were designed to be used).

4.3. The Time is Right…

The down-side of the distributed, agent-based approach described in Section 4.1 is a reduction in efficiency of operation – it is always the case that a custom-built application can be designed in a way that will out-perform the equivalent that is built from distributed components. Throughout the history of computer systems, when a new application area is initially developed, flexibility typically has to be compromised in order to achieve a certain level of performance. As time goes on, the technologies involved

mature – hardware becomes more powerful and software becomes more sophisticated. These technology improvements mean that it gradually becomes possible to introduce the extra flexibility and functionality associated with modularity without compromising performance.

We believe that now is the right time to extend the concepts of modularity right down to the intelligent sensor level. Technology has reached a point where the overhead associated with this approach is acceptable, and we also have well-enough developed software techniques [jini][upnp] to build intelligent sensor networks in this way.

4.4. Smart Sensors in an Auto-ID Framework

The smart sensor network we imagine would be a rich source of additional information for the applications the Auto-ID Centre envisages. In addition to information about the identity of objects, smart sensors would generate other data such as location, temperature and shock/vibration, and the smart sensor network would provide a means to communicate this. The Auto-ID systems would then either collate and filter this information [oats2002] before passing it up to higher-level business information systems or they would process the information directly and use it to control the environment [mcfarlane2002]. A smart sensor network clearly has features that mean that it can provide a valuable supporting role in an Auto-ID system. Conversely, the Auto-ID system can support the smart sensor network. For example, information about individual smart sensor devices, their abilities and their configurations could be stored as PML so that it is readily accessible not only to the corresponding device, but also to other devices in the network. In this way, the 'driver' information needed for plug'n'play operation can be easily maintained and distributed.

One of the activities within the Auto-ID Centre is the development of universal protocols suitable for RFID tag communication. Since RFID is a key smart sensor technology, it seems sensible to make sure that the Auto-ID centre protocols are extensible to the more generic and sophisticated ideas presented in this paper. This would provide a common platform upon which a huge variety of applications can be built.

5. PROBLEM AREAS

In the previous sections we demonstrated the applications and potential benefits a universal smart sensor network could bring to all parties involved in developing, producing, transporting and selling products. However, in order to realise the vision, many problems still need to be solved. Some of these problems are of a technical nature, but many are not. Table 1 visualises what we have identified as the four main problem areas – the four cornerstones as it were – in building a universal smart sensor network. The following sections will look into these four areas in more detail.

Table 1: The four cornerstones for building a universal smart sensor network

UNIVERSAL SMART SENSOR NETWORK			
Standards	**Collaboration**	**Software Frameworks**	**Hardware**
– Object Identification – Object Description – e-Business Processes – Software Interfaces and Protocols	– Sharing of Business Data – Deployment – Changes of Ownership	– Adaptive – Device/Service Discovery – Scalability – Security	– Ubiquitous Sensors – Cheap Components – Automatic Sensor Calibration – Length of Life

The collaboration and standardisation issues described below take often more time to solve than the technical issues. They can only be solved through business decisions, which will only be taken in favour of a universal smart sensor network if the right business cases can be proven.

5.1. Standards

By definition, a universal smart sensor network can only come into life if it is a network as common as the Internet is today – we want to create the "Internet of Things". In order to achieve this, global standards need to be developed and applied in the following areas:

Object identification
A global object identification scheme is important in order to track single items and to avoid object reference errors that can occur when object identifiers need to be converted or mapped. The electronic product code (EPC™) as defined by the Auto-ID Centre satisfies this requirement.

Object description
Perhaps more interesting than the identity of objects is the ability to exchange information about those objects. This requires standards relating to data formats and semantics, e.g. the meaning of a specific XML tag should be the same for all participants in a supply chain. Many organisations have defined such standards for their domain or are working on defining new ones. Depending on the organisation, these standards can be industry-specific or cross-industry [rosettanet, exbml]. The physical markup language (PML) as proposed by the Auto-ID Centre [brock2001b] is another format for storing information about objects. PML is divided into core elements, e.g. time and location, and PML extensions [floerkemeier2002]. This division is in our opinion one of the strongest points of PML since it enables the reuse of the work of other standard bodies. For example, ebXML [ebxml] tags or tags defined by RosettaNet [rosettanet] can be used within PML.

e-Business processes
The business processes of partners must be aligned in order to know what data needs to be exchanged and when. The work of the different standards organisations like RosettaNet and ebXML can be reused for this.

Software interfaces and protocols
To enable the ubiquitous deployment envisaged, the software interfaces and protocols between single software components need to be standardised as well. It is important that only the interfaces and protocols are standardised, not the components themselves. This ensures that technology vendors have the flexibility to add value in their own ways, but without compromising interoperability. We imagine that vendors will need the freedom to adopt different (potentially proprietary) implementation approaches, or perhaps to provide specific, additional functionality for specific application areas; this must be supported without compromising the interoperability of the sensor devices.

While standards are very important, the authors strongly believe that just waiting for the right standards to be developed is not appropriate. This would take too much time and there would be a high probability that the standards wouldn't actually be suitable for daily use in the field. The standards have to be developed in connection with practical field tests to make sure that they are actually useful and relevant, and to propagate them throughout the industry. The difficulty though will be to walk the fine line between implementing particular scenarios for specific field tests and producing generalised standards that are applicable for other scenarios as well. The proposed solution must be flexible and powerful enough that it can address the needs of the entire community, but at the same time it must be lean and tight enough to provide a practical solution in each of the specific applications needed by individuals within the community.

5.2. Collaboration

A universal smart sensor network can provide a lot of data to enterprises about what is really happening in their supply chains and how they are executing. It has been shown [joshi2000] that to have complete visibility throughout the supply chain reduces the inventory and backlog costs dramatically throughout the chain. However, such visibility can only be achieved when companies are willing to share internal data with business partners and unfortunately competitive concerns limit the willingness of companies to do so. One way of alleviating these concerns is to use software agents to communicate information instead of sharing the raw data [huhns2002]. The software agents act as gatekeepers that can answer questions about availability of a product or a resource in a certain time frame without giving away more data than absolutely necessary, such as exactly how much stock of a given product is being held. In a further refinement, the agents could also negotiate price and delivery times.

Note that even if the costs of the whole supply chain are reduced, a company will only be willing to participate if there is a direct financial benefit for itself, for example by being able to keep lower inventory levels or perhaps by being explicitly reimbursed in some kind of costs savings distribution scheme. A similar issue relates to deployment. Deploying the necessary hardware and software infrastructure has some considerable cost. The standardisation work will bring down the cost of the components in the infrastructure, but it will not be zero. A cost-sharing scheme where the cost of deployment is split among the different participants according to the individual benefits would be helpful. Take for example the tagging of individual products. Ideally, products are tagged at the point of manufacture. In many cases the financial benefits that tagging can bring for the manufacturer alone justify the costs. However, should the manufacturer bear the whole cost associated with source level tagging, when many of the benefits of tagged products will go to the retailer, not the manufacturer? Or should these costs be shared more evenly?

Not only will many of the smart sensors themselves move along the supply chain during the life of a product, but also the information associated with them and with the associated products will move between computer systems belonging to different corporate entities. There are subtle questions regarding the ownership of the sensors and of the information during the lifetime of a product – at what point, for example, does the information leave the domain of the manufacturer, and come under control of the retailer? Is the information partitioned so that there is a joint ownership of some kind? How is access for reading, overwriting and deleting existing information controlled? How is support for adding new information provided? Who is responsible for ensuring data integrity, in the face of both accidental and malicious events? These are all interesting technical and political questions, which need to be addressed through a combination of research, experimentation and dialog between interested parties.

If solutions for these collaboration issues cannot be found, a truly universal smart sensor network will not come into existence. However, the technology and standards will still allow for local smart sensor networks deployed within single corporations or specific consortia. We believe that there are also significant benefits for such local smart sensor networks.

5.3. Software Frameworks

To support a universal smart sensor network requires a software infrastructure – based on standard interfaces and protocols – with the following properties:

Openness
Heterogeneous environments, both regarding software as well as hardware, are a fact of life. Standardisation will never and should never go so far as to define one given infrastructure for everybody. We therefore prefer to use the term framework instead of infrastructure, because it implies that some adaptations to specific situations are possible and perhaps even necessary. It must be possible for different vendors to add or integrate new components easily.

Scalability
The large number of devices and sensors integrated into the network will generate a huge amount of data that the network must be capable of handling. There are a number of ways in which this may be achieved: partitioning the network, network hierarchies, data filtering and aggregation at several layers within the network hierarchy, local control decisions etc.

Security
Needless to say that security is a crucial issue in any IT-related infrastructure. Security in the context of a smart sensor network takes many forms, such as privacy, authorisation and authentication, integrity and non-repudiation. The sensors that generate data and the actuators that use data to control the environment will be highly distributed, creating new challenges for the design and implementation of security policies. An insecure solution can of course lead to financial losses or even worse, breaches of safety.

Device and service discovery
A network of this scale can only work if it is possible to plug a new device into the network without the need for active configuration of either the device or other components that would like to use the services this device provides. It is desirable that the roles of and the relationships between the devices and the business processes and enterprise systems are discovered automatically [balazinska2002]. At the same time, the enterprise systems must be able to make use of new service autonomously. Devices can also disappear from the network at unpredictable times. This must be detected and appropriate action needs to be taken: rerouting service requests to other sensors or in same cases alerting a technician about the failure of a critical device, so that the device can be fixed.

Adaptation
We have already mentioned two facets of adaptation, namely: adapting the framework to specific scenarios and automatic adaptation as devices appear on or disappear from the network. A third aspect involves adapting the behaviour of the network over time. Information about what actually happens in the network and how resources are utilised can be used for more efficient resource allocation and configuration and even change the behaviour of some of the components and agents in the network.

5.4. Hardware

As the name implies, a smart sensor network consists of many networked devices. Advances in hardware technology are necessary to realise the vision for a truly universal, robust and reliable network. Some of the issues are:

Price
Large-scale deployment will only be possible if the individual components are very cheap. The Auto-ID Centre is working on developing the technology for very cheap tags and readers [sarma2001, reynolds2002], but this needs to be extended to other forms of smart sensors.

Length of life
How long can a sensor keep operating reliably? This mainly depends on how the sensor is powered, since most sensors will need some form of power at least for performing continuous sensing (of temperature for example), if not for communicating the sensed data. Battery and low power technologies then are the determining factors how long an autonomous sensor can be deployed. A lot of research is currently being done in this area, for example within the “Smart-Its” research program [smart-its].

Sensor calibration
Sensors deployed in the field tend to become uncalibrated over time, resulting in inaccurate data. There are several ways this problem can be overcome. One can deploy a multitude of sensors sensing the same parameter at a specific location. Voting algorithms can be used to determine the most likely value; this value is then propagated and used to recalibrate the sensors. This can be done on a continuous basis. Another solution is to periodically replace or recalibrate the sensors although the effort this takes could be prohibitive. One interesting way of doing recalibration is to use a robot that periodically visits the sensors [lamarca2002]. This is only feasible if the locations have to be visited regularly anyway, for example by a maintenance robot that could easily take on this additional task of recalibration.

Ad-hoc networking
A related problem is how easily we can deploy sensors and connect them to power and the network without effort. Wireless communication technologies are often too power-consuming. An alternative could be the use of embedded conductive materials in walls and other surfaces. A sensor can then be pinned to the wall and will immediately have power and a network connection [laerhoven2002, lifton2002].

6. CONCLUSION

In this paper, we have outlined our vision of a universal smart sensor network and showed the applications such a network would enable. We also highlighted some of the underlying technologies upon which such a network could be built, and given an overview of the issues where further work is necessary in order to realise the vision. Now is the time to start developing the architecture and protocols for smart sensor networks. As RFID is one of the key enabling technologies, it is sensible to dovetail with protocols for RFID tags in order to provide a common basis on top of which applications can be built and which can be extended to cover other types of smart sensor technologies.

We believe that the way to move forward is to develop and deploy different types of smart sensor technology, and implement applications on top of this framework. Only by doing so can we learn about the practical issues, the benefits and of course, the pitfalls. This will ensure that the standards developed are of practical relevance and will become as ubiquitous as the Internet is today, thereby preventing divergence of the field into several incompatible solutions.

7. REFERENCES

[aim]

1. **"Bar Code Technology".**
   Available at http://www.aimglobal.org/technologies/barcode/

[alexander2002a]

2. **K. Alexander, G. Birkhofer, K. Gramling, H. Kleinberger, S. Leng, D. Moogimane & M. Woods, "Focus on Retail: Applying Auto-ID to Improve Product Availability at the Retail Shelf".**
   Auto-ID Centre business case IBM-AUTOID-BC-001, 1/6/2002.
   Available at http://www.autoidcenter.co.uk/research/IBM-AUTOID-BC-001.pdf

[alexander2002b]

3. **K. Alexander, T. Gilliam, K. Gramling, M. Kindy, D. Moogimane, M. Schultz & M. Woods, "Focus on the Supply Chain: Applying Auto-ID within the Distribution Center".**
   Auto-ID Centre business case IBM-AUTOID-BC-002, 1/6/2002.
   Available at http://www.autoidcenter.co.uk/research/IBM-AUTOID-BC-002.pdf

[badrinath1996]

4. **B. Badrinath, T. Imieliski, R. Frenkiel & D. Goodman, "NIMBLE: Many-Time, Many-Where Communication Support for Information Systems in Highly Mobile and Wireless Environments".**
   Available at http://www.cs.rutgers.edu/~badri/dataman/nimble/

[balazinska2002]

5. **M. Balazinska, H. Balakrishnan & D. Karger, "INS/Twine: A Scalable Peer-to-Peer Architecture for Intentional Resource Discovery".**
   Proceedings of First International Conference on Pervasive Computing, Pervasive 2002, Zürich, Switzerland, August 26–28, 2002. Springer Verlag Heidelberg, Lecture Notes in Computer Science Vol. 2414, p. 195 ff.

[barnes1999]

6. **D. Barnes, J. Bradshaw, L. Day, T. Schott & R. Wilson, "Two Dimensional Bar Coding".**
   Tech 621 course notes, Purdue University, Spring 1999.
   Available at http://www.tech.purdue.edu/graduate/courses/Tech621aw/2Dbarcodes.PDF

[booch1994]

7. **G.Booch,"Object-Oriented Analysis and Design with Applications".**
   Addison Wesley Professional, second edition, 1994. ISBN: 0-8053-5340-2

[brock2001a]

8. **D. Brock, "The Electronic Product Code (EPC™) – A Naming Scheme For Physical Objects".**
   Auto-ID Centre white paper MIT-AUTOID-WH-002, 1/1/2001.
   Available at http://www.autoidcenter.co.uk/research/MIT-AUTOID-WH-002.pdf

[brock2001b]

9. **D. Brock, T. Milne, Y. Kang & B. Lewis, "The Physical Markup Language".**
   Auto-ID Centre white paper MIT-AUTOID-WH-005, 1/6/2001.
   Available at http://www.autoidcenter.co.uk/research/MIT-AUTOID-WH-005.pdf

[bussmann1998]
10. **S. Bussmann, "An agent-oriented architecture for holonic manufacturing control".** Proccedings of the 1st international workshop on Intelligent Manufacturing Systems, IMS-Europe 1998, EPFL, Lausanne, Switzerland

[chaxel1997]
11. **F. Chaxel, E. Bajic & J. Richard, "Mobile DataBase Nodes for Manufacturing Information Management: a STEP based Approach".** International Journal of Advanced Manufacturing Technology, Vol 13, 1997, pp 125–133

[coulouris1988]
12. **G. Coulouris & J. Dollimore, "Distributed Systems: Concepts and Design".** Addison-Wesley, 1988

[ebxml]
13. **Electronic Business using eXtensible Markup Language website.** http://www.ebxml.org/

[finkenzeller1999]
14. **K. Finkenzeller, "RFID Handbook".** 1st edition, Wiley & Sons LTD, 1999, ISBN 0-471-98851-0

[fleisch2001]
15. **E. Fleisch, "Business perspectives on Ubiquitous Computing".** M-Lab Working Paper No. 4, Version 1.0, 30/11/01. Available at http://www.inf.ethz.ch/vs/m-lab/WP4.pdf

[floerkemeier2002]
16. **C. Floerkemeier & R. Koh, "Physical Mark-Up Language Update".** Auto-ID Center 2002. Available at http://www.autoidcenter.org/research/MIT-AUTOID-TM-006.pdf

[gorday2001]
17. **P. Gorday (Motorola, 8000 W. Sunrise Blvd., M/S 2141, Plantation, FL 33322), J. Gutierrez (Eaton) & P. Jamieson (Philips), "IEEE 802.15.4 Overview".** Doc. IEEE 802.15-01/358r0, 12 November, 2001. http://grouper.ieee.org/groups/802/15/pub/2001/Nov01/01509r0P802-15_TG4-Overview.ppt

[huhns2002]
18. **M. Huhns, L. Stephens & N. Ivezic, "Automating supply-chain management".** In proceedings of the First International Joint Conference on Autonomous Agents and Multi-Agent Systems (AAMAS 2002), July 15–19, 2002, Bologna, Italy, pp. 1017–1024

[irda]
19. **Infrared Data Association Website** http://www.irda.org/

[jini]
20. **Jini Website** http://www.jini.org/

[joshi2000]

21. **Y. Joshi, "Information Visibility and its Effect on Supply Chain Dynamics".**
Master Thesis, Department of Mechanical Engineering, Massachusetts Institute of Technology 2000.
Available at http://www.autoidcenter.org/research/YVJ-Thesis.pdf

[kambil2002]

22. **A. Kambil & J. Brooks, "Auto-ID Across the Value Chain: From Dramatic Potential to Greater Efficiency & Profit".**
Auto-ID Centre business case ACN-AUTOID-BC-001, 1/6/2002.
Available at http://www.autoidcenter.co.uk/research/ACN-AUTOID-BC-001.pdf

[kubach2001]

23. **U. Kubach & K. Rothermel, "Exploiting Location Information for Infestation-Based Hoarding".**
In Proceedings of the Seventh ACM SIGMOBILE Annual International Conference on Mobile Computing and Networking (MobiCom 2001), pp. 15–27, Rome, Italy, July 2001

[langheinrich2002]

24. **M. Langheinrich, "Privacy Invasions in Ubiquitous Computing".**
Workshop on Socially-informed Design of Privacy-enhancing Solutions in Ubiquitous Computing. 4th International Conference on Ubiquitous Computing, UbiComp 2002, September 2002.
Available at http://www.inf.ethz.ch/vs/publ/papers/uc2002-pws.pdf

[laerhoven2002]

25. **Pin &Play: K. Laerhoven, A. Schmidt & H. Gellersen (Lancaster University), "Networking Objects through Pins (tech note)".**
Proceedings of Fourth International Conference on Ubiquitous Computing, UbiComp 2002, Göteborg, Sweden, September 29 – October 1, 2002. Springer Verlag Heidelberg, Lecture Notes in Computer Science Vol.2498, p.219 ff.

[lamarca2002]

26. **A. LaMarca, W. Brunette, D. Koizumi, M. Lease, S. Sigurdsson, K. Sikorski, D. Fox & G. Borriello, "Making Sensor Networks Practical with Robots".**
Proceedings of First International Conference on Pervasive Computing, Pervasive 2002, Zürich, Switzerland, August 26–28, 2002. Springer Verlag Heidelberg, Lecture Notes in Computer Science Vol. 2414, p. 152 ff.

[leick1995]

27. **A. Leick, "GPS Satellite Surveying".**
2nd Edition, John Wiley & Sons, Inc (New York), 1995. ISBN: 0-471-30626-6

[lifton2002]

28. **J. Lifton, D. Seetharam, M. Broxton & J. Paradiso, "Pushpin Computing System Overview: A Platform for Distributed, Embedded, Ubiquitous Sensor Networks".**
Proceedings of First International Conference on Pervasive Computing, Pervasive 2002, Zürich, Switzerland, August 26–28, 2002. Springer Verlag Heidelberg, Lecture Notes in Computer Science Vol. 2414, p. 139 ff.

[luck2001]
29. **M. Luck, V. Marik, O. Stepankova & R. Trappl (Eds.), "Multi-agent systems and applications".** 9th ECCAI Advanced Course, ACAL 2002, Prague, July 2001. Springer-Verlag Lecture Notes series, ISBN 3-540-42312-5.

[mcfarlane2000]
30. **D. McFarlane & S. Bussman, "Developments in Holonic Production Planning and Control".** Int. Journal of Production Planning and Control, Vol. 11, No 6, 2000, pp. 522–536

[mcfarlane2002]
31. **D. McFarlane, "Auto-ID Based Control – An Overview".** Auto-ID Centre white paper CAM-AUTOID-WH-004, 1/2/2002. Available at http://www.autoidcenter.co.uk/research/CAM-AUTOID-WH-004.pdf

[miller 2001]
32. **B. Miller & C. Bisdikian, "Bluetooth Revealed".** Prentice Hall PTR, 2001, ISBN 0-13-090294-2

[oats2002]
33. **Oat Systems and MIT Auto-ID Center, "The Savant™ – Version 0.1 (Alpha)".** Auto-ID Centre technical memo MIT-AUTOID-TM003, 1/2/2002. Available at http://www.autoidcenter.co.uk/research/MIT-AUTOID-TM003.pdf

[oliphant2000]
34. **M. Oliphant, "Radio interfaces make the difference in 3G cellular systems".** IEEE Spectrum, October 2000. Available at http://www.spectrum.ieee.org/publicfeature/oct00/cell.html

[philips2002]
35. **"Philips and Sony Announce Strategic Cooperation to Define Next Generation Near Field Radio-Frequency Communications".** Philips press release, 5/9/2002. Available at http://www.semiconductors.philips.com/news/content/file_878.html

[pappu2002]
36. **R. Pappu, B. Rect, J. Taylor & N. Gershenfeld, "Physical One-Way Functions".** Science, Vol 297 (5539) 2026, September 2002

[reynolds2002]
37. **M. Reynolds, J. Richards, S. Pathare, H. Tsai, Y. Maguire, R. Post, R. Pappu & B. Schoner, "Multi-Band, Low-Cost EPC™ Tag Reader".** Auto-ID Centre white paper MIT-AUTOID-WH-012, 1/6/2002. Available at http://www.autoidcenter.co.uk/research/MIT-AUTOID-WH-012.pdf

[riezenman2002]
38. **M. Riezenman, "The ABCs of IEEE 802.11".** http://www.spectrum.ieee.org/WEBONLY/resource/sep02/802ABCs.html

[rosettanet]
39. **Rosetta net website.** http://www.rosettanet.org/

[sarma2001]

40. S. Sarma, "Towards the 5¢ Tag".
Auto-ID Centre white paper MIT-AUTOID-WH-006, 1/11/2001.
Available at http://www.autoidcenter.co.uk/research/MIT-AUTOID-WH-006.pdf

[sarma2002]

41. S. Sarma, S. Weis & D. Engels, "Low Cost RFID and the Electronic Product Code".
Proceedings of the Workshop on Cryptographic Hardware and Embedded Systems 2002 (CHES 2002), August 2002.

[smart-its]

42. Smart-Its research project, European Disappearing Computer Initiative.
http://www.disappearing-computer.net/projects/SMARTITS.html

[thede2001]

43 A. Thede, A. Schmidt & C. Merz, "Integration of goods delivery supervision into E-Commerce supply chain".
Second International Workshop on Electronic Commerce (WELCOM'01).
November 16–17, 2001. Heidelberg, Germany.
Available at http://www.comp.lancs.ac.uk/~albrecht/pubs/pdf/thede_welcom_2001.pdf

[upnp]

44. Universal Plug and Play Forum.
http://www.upnp.org/

[wooldridge1997]

45. M. Wooldridge. "Agent-based Software Engineering".
In IEE Proceedings on Software Engineering, 144(1), pages 26–37, February 1997.
Available at http://www.csc.liv.ac.uk/~mjw/pubs/iee-se.pdf

[ye1998]

46. T. Ye, H.-A. Jacobsen & R. Katz, "Mobile Awareness in a Wide Area Wireless Network of Info-Stations".
Proceedings of the Fourth Annual ACM/IEEE International Conference on Mobile Computing and Networking (MobiCom '98)

Designed and produced by Foxner. www.foxner.com