WHITE PAPER SERIES / EDITION 1

AUTO–ID LABS

BUSINESS PROCESSES & APPLICATIONS
SOFTWARE & NETWORK
HARDWARE

AUTOIDLABS-WP-SWNET-015

Auto ID — 21st Century Supply Chain Technology

Peter H. Cole, Daniel W. Engels

School of Electrical and Electronic Engineering
The University of Adelaide, Adelaide SA 5005
cole@eleceng.adelaide.edu.au

Massachusetts Institute of Technology
77 Massachusetts Avenue, Room 3-449, Cambridge, MA 02139 USA
dragon@lcs.mit.edu

www.autoidlabs.org

Abstract

We describe the Networked Physical World system, a combination of RFID technology and a ubiquitous computing system that integrates the physical world with the virtual world. The instrumentation of non-electronic devices, such as cans of soda and boxes of detergent, with inexpensive, low functionality computing devices enables the virtual world to sense and identify physical objects. The ability to accurately and definitively identify physical objects without human intervention is paramount to many applications including home automation and supply chain management. However, a system that networks all physical objects requires a careful, deliberate system design to support the huge scale of a quadrillion node network. A prototype of the Networked Physical World system has been implemented and we are evaluating its performance and capabilities within a large-scale real-world supply chain application, the Field Trial. We describe the building block system components of the Networked Physical World system and present some preliminary results from the Field Trial.

1. Introduction

The Auto-ID Center is an industry sponsored non-profit organization, with laboratories in four continents, chartered to design the open standards-based system that connects all physical objects to the global Internet. The resulting system, the Networked Physical World system, creates a "Smart World," that is, an intelligent infrastructure linking objects, information, and people through computer networks. This new infrastructure enables object-centric computing that will allow universal coordination of physical resources through remote monitoring and control by both humans and machines. The Auto-ID Center's objective is to create open standards, protocols, and languages to facilitate worldwide adoption of this Networked Physical World system — forming the basis for the new "Internet of Things".

The availability of a low-cost Networked Physical World system connecting physical objects to the Internet will have an immediate and profound impact on supply chain management – and bring the vision of the Auto-ID Center to a near-term reality. It is our belief that such a low-cost system is possible, and the Auto ID Center is undertaking research and development activities to achieve this goal. The implementation and deployment of the Networked Physical World system prototype are the first steps of these activities.

The Networked Physical World system is an object-centric system explicitly designed to allow tasks to be performed on behalf of physical objects. That is, physical objects are at the center of the Networked Physical World system. This is markedly different from the human-centric approach taken by traditional ubiquitous and pervasive computing projects such as the Cooltown project [Barton-2001], the Portolano project [Esler-1999], and the Oxygen project [MIT-2000]. Object-centric computing moves beyond human-centered tasks that assist humans in their daily activities and enables complex tasks to be performed on behalf of objects, eliminating the need for humans to perform any part of these tasks.

The core components of the Networked Physical World system include a universal identification scheme, a generic redirection service, a local data collection and control system, a generic object description language, and low-cost electronic tags used to connect physical objects to the Internet. These components and their prototype reference implementations form the basis for a set of open standards, protocols, and languages designed to connect physical objects to the global Internet.

We also describe a large-scale, real-world Field Trial that fully employs the systems, standards, and languages developed at the Auto-ID Center. This Field Trial is a multi-city, multi-company trial of the system in a supply chain application, in which thousands of items, from pallets to consumer products, are tagged and tracked, in real-time, from manufacturing facilities through distribution centers to and within the retail store.

Section 2 provides an overview of the Networked Physical World system and its primary components. The general properties of RFID tags, the wireless connection between the physical and virtual worlds, and the Center's research in this area are described Section 3. Section 4 describes the Savant data collection and automatic control system that forms the ubiquitous data collection and localized control building block of the Networked Physical World system. Section 5 describes the Electronic Product Code, a unique identification scheme for physical objects. The Object Name Service, a redirection service similar to DNS, is described in Section 6. Section 7 describes the Physical Markup Language, a language for describing physical objects. Section 8 describes a large scale Field Trial of the Networked Physical World system, and Section 9 summarizes the system.

2. System Overview

The Networked Physical World system is an intelligent ubiquitous infrastructure that automatically and seamlessly links physical objects to the global Internet. This system networks physical objects without human intervention or manipulation by automated machines. This is accomplished by integrating an electronic Radio Frequency Identification (RFID) transponder, or tag, into each object. The system networks objects seamlessly by communicating with these tags at suitably placed locations e.g., portals, at mobile locations, e.g. through handheld devices, and potentially, eventually for some tags, continuously throughout the environment. A network of tag readers and local data collection and control systems, called Savants, is used to automatically communicate with the tags of physical objects, and to automate control applications.

The ubiquitous nature of the Networked Physical World system requires that the system be inexpensive to implement relative to the benefits achieved by applications that utilize the system, such as supply chain management. The extreme low cost of tags and readers required to actually implement the system has been the overriding constraint on the design of the Networked Physical World system.

The cost of tags for quadrillions of objects is the dominant cost of the system. Consequently, the tag cost, and therefore its functionality, must be minimized. With few exceptions, the resulting cheap tag stores only a unique identifier, the Electronic Product Code (EPC), for the object.

The unique object identifier is global in scope and acts as a pointer to information stored about the object somewhere over the network. A redirection service, the Object Name Service (ONS), is used in conjunction with the EPC to identify the location(s) of information and related services for a particular object. The ONS allows for the location of locally available information as well as globally available information.

The information must be stored in a standard language to enable true automation. The Networked Physical World system utilizes the Physical Markup Language (PML), an XML based language, to describe physical objects and their properties.

Whenever possible, the Networked Physical World system adopts and uses standards either developed by or in cooperation with governing bodies, such as the Internet Engineering Task Force (IETF), the Institute of Electrical and Electronic Engineers (IEEE), the Uniform Code Council (UCC) and EAN International, the American National Standards Institute (ANSI), and the International Standards Organization (ISO), as well as commercial consortium and industry groups.

The Networked Physical World intelligent infrastructure has five primary components, some physical, some, conceptional: (1) electronic RFID tags, (2) the Savant systems, (3) Electronic Product Code (EPC), (4) Object Name Service (ONS), and (5) Physical Markup Language (PML). We provide a more detailed description of these components in the following sections.

3. RFID Tags

Electronic tagging refers to a family of technologies that transfer data wirelessly between tagged objects and electronic tag readers. Traditional bar codes may be considered electronic tags under this definition since the printed bars reflect light and communicate their data to the laser scanner. Electronic Article Surveillance (EAS) labels also communicate information wirelessly, but only a single bit of data — whether or not an item is in the field.

Radio Frequency Identification (RFID) tags, often used in „smart cards," have a small radio antenna that transmits information over a short range to an RFID tag reader [Finkenzeller-1999]. RFID technology may use both powered and non-powered means to activate the electronic tags. Powered devices use batteries to actively transmit data from the tags to more distant readers. Electronic highway toll systems are good examples of active RFID tags. Passive RFID devices literally harvest energy from the electromagnetic field of an active reader to both power the tag and transmit the data.

Historically, many physical principles for RFID data encoding and extraction have been used, but the focus now is on monolithic microelectronic technology labels interrogated by the process of RF backscatter. [Cole-1996]. Electromagnetic waves of frequencies ranging from 100 kHz to 3 GHz are used. Such waves have varying degrees of success in penetrating surrounding materials. The operation of RFID labels is significantly restricted by electromagnetic compatibility considerations. In the main, RFID is expected to maintain an uneasy coexistence with many other RF systems in ISM bands.

In the most popular technology, the tags are passive, i.e. they contain no energy source, and in consequence the ranges are short - only up to a few metres, the latter being achievable only at optimum frequencies and in regulatory environments sympathetic to RFID. [Cole-1995].

These passive systems are very well suited for use in the Networked Physical World via the EPC concept, but our vision extends beyond passive tags to other classes of tags with generally and sometimes substantially enhanced functionality.

Figure 1 shows in simple form some features of RFID systems used in the Networked Physical World application, i.e. the plurality of labels simultaneously present in the interrogation field, and the weak nature of the reply in comparison with the interrogation signal. This latter feature results from the accumulation of the electromagnetic propagation loss in both directions of the electromagnetic link, together with the absence of gain in power in the label (in fact there is a loss), and the fact that the energy in the interrogation signal has to be divided among the many bits of the tag reply. These constraints are not shared by many applications of electronic communication with which we are familiar, and care is needed before applying intuition gained in those other contexts to RFID applications.

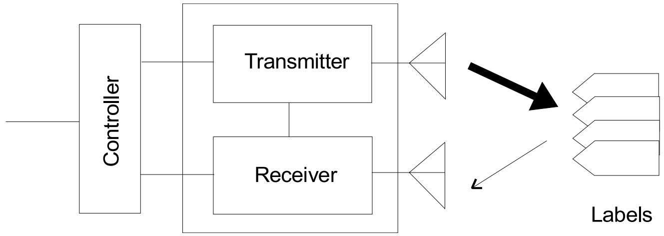

Fig. 1

The functionality of a typical RFID tag is contained within a small microchip connected to an antenna (and a battery in the case of powered tags). The more functionality that a tag implements, the larger and more expensive is the microchip. The Auto-ID Center has specified the air interface for a cheap RFID tag that requires a minimum of functionality implemented on the microchip. This functionality consists solely of a means to efficiently read a single unique identifier stored on the tag. Some of our sponsoring companies have designed and are manufacturing cheap passive tags compliant with the Center's specification for use within our Field Trial. It is expected that these tags will cost 5¢ or less in volumes of several billion tags [Sarma-2001].

Due to the cost constraints of tagging a quadrillion objects, it is expected that the majority of tags integrated with physical objects will be passive tags.

Electronic tags, when coupled to a reader network, form the link between physical objects and the virtual world in our system. A ubiquitous reader network allows continuous tracking and identification of physical objects. Reader arrays have been fabricated and integrated in floor tiles, carpeting, shelf structures, cabinets and appliances. Similarly to cellular phone grids, the reader network may provide seamless and continuous communication to tagged objects. A data collection and control system must support the reader network to enable efficient use of the continuous, or at least very frequent, object communications. Additionally, in order to access and identify these objects, we need a means to uniquely name them and store and retrieve information about them.

4. Savant System

A Savant system is an event router and local control system that performs operations such as data capture, data monitoring and filtering, and data transmission [Auto-ID Center-2001a]. Networked Savant systems form a framework to manage and react to the EPC values communicated to the tag readers.

In its most general form, the Savant system acts as a gateway to local networks, devices, data storage, and inference engines—generating and receiving requests for information. Subsets of this general definition may include Savants without any local network or device, simply database query engines. Others include only physical devices, such as sensors, which return measured data, or controllable mechanisms that respond to commands, such as a thermostat controller. Some Savants generate requests for data autonomously, while others simply respond to received commands.

The Savant systems are deployed in a hierarchical, distributed framework. This topology enables the Savant framework to handle large volumes of communicated EPCs from networked objects. This flexible framework is well suited for business applications, such as real-time inventory management, out-of-stock identification or notification, and theft prediction. The Savant system can interface with corporate Enterprise Resource Planning (ERP) systems, thereby providing accurate real-time item level information to these systems.

At the heart of a Savant node is the task manager, which schedules, executes and manages multiple processes. The task manager schedules both recurring and preemptive tasks and facilitates the recovery of task state from system shutdown and restart. Common tasks include data migration, ONS and PML communication, data caching, and application specific processes such as theft prediction.

Communication with the Savant, including requests for information, transmission of data and the reception of new task processes, are handled through the SOAP interface. The SOAP interface allows access to the services offered by the Savant, such as query processing and task management. The SOAP interface may also exchange information with third party database systems.

A class server provides platform-independent Java classes to the Savant. The Java Virtual Machine in the Savant can load classes remotely, on-demand, and from the class server.

The query processor is a one-time task that interprets remote queries to the Savant. The task manager spawns a query processor for each query and returns the results to the remote caller. Queries can access both local resources on the savant, as well as remote services such as ONS and PML data files. Queries are written in Distributed Python, a pure-Java implementation of Python that provides libraries for process management and data communication.

A series of software components and utilities has been developed to increase the functionality and practicality of the Savant implementation. These include

- an EPC event filter that more accurately reflects the state of the environment,
- a remote administration and management system for the Savant architecture, and
- prototypes of and concepts for security, authentication and access control systems.

5. Electronic Product Code

The Electronic Product Code (EPC) is a globally unique identification scheme designed to uniquely identify all physical objects and aggregations of objects [Brock-2001a][Brock-2001c]. The EPC code is sufficiently large to enumerate all objects and to accommodate all current and future naming methods. It provides for industry coding standards, such as those from the Uniform Code Council (UCC) and the EAN International. These standards include the original Uniform Product Code (UPC), as well as other numbering schemes, such as the Shipping Container Code (SCC-14) and the Serial Shipping Container Code (SSCC-18) [UCC-1986][UCCa][UCC-b][UCC-c]. The EPC was intended, as much as possible, to be universally and globally accepted.
Since the EPC is used primarily to link physical objects to the network, it was designed to serve as an efficient information reference. Finally, the code is extensible, allowing future expansion in both size and design.

The EPC scheme consists of four distinct, hierarchical partitions: version number, domain manager number, object class code, and serial number. All EPC codings contain these four partitions. The first partition, the version number, contains information on the length and structure of the code being used, and the three remaining partitions contain the actual unique identifier for the object. The EPC scheme does not contain any error detection or correction coding since the correct coding to use depends upon how a specific EPC is communicated, but such detection or correction is implemented readers.

The 96-bit EPC, EPC-96 Type I, is shown in Figure 2. The EPC version number spans 8 bits, theoretically allowing up to 256 different 96-bit versions, with the possibility that smaller and larger EPC versions with different partitioning, still recognisable for the header, may be devised.

The EPC manager partition spans a 28-bit section, encoding a maximum of 228 = 268,435,456, or approximately 268 million, manufacturers.

The next partition, object class code, occupies the next 24-bits. The object class may be considered the product stock-keeping unit (SKU). It may also be used for lot number, or any other object-grouping scheme developed by the EPC manager. Since each manufacturer is allowed more than 16 million object types, this partition could encode all of the current UPC SKUs, as well as many other object classes. This is important as applications expand beyond current retail applications into general supply chain and manufacturing.

Bits 0, …, 7	EPC version number
Bits 8, …, 35	Domain manager number
Bits 36, …, 59	Object class number
Bits 60, …, 95	Serial number

Fig. 2: The EPC-96 Type I encoding.

The EPC's final partition encodes an object identification serial number, unique for all objects of that domain and class. For all objects of a similar type, the EPC serial number provides 36-bits, or 236 = 68,719,476,736, unique identifiers. Together with the product code, this provides each manufacturer with 1.1x1018 unique item numbers – currently beyond the range of all identified roducts in the world.

6. Object Name Service

The Object Name Service (ONS) is the "glue" that links the Electronic Product Code with its associated Physical Markup Language (PML) data file(s). The ONS is a system designed to automatically locate networked information and services for tagged objects. More specifically, the ONS is an automated networking service, which, when given an EPC number, returns a set of host addresses on which the corresponding PML files or services are located [Auto-ID Center-2002b].

The ONS architecture is a distributed framework consisting of (1) ONS servers, (2) Mapping information and (3) an ONS resolver.

The ONS server responds to queries requesting the IP addresses of PML servers for a specific EPC or contiguous range of EPCs. Servers having authoritative information regarding certain EPCs return immediately, while others return a pointer to an ONS server that will lead to the authoritative ONS server for that EPC. Mapping information is hierarchically distributed across the ONS servers for manageability. Mapping information is stored in records that express the mapping between a "range" of EPCs and a PML server.

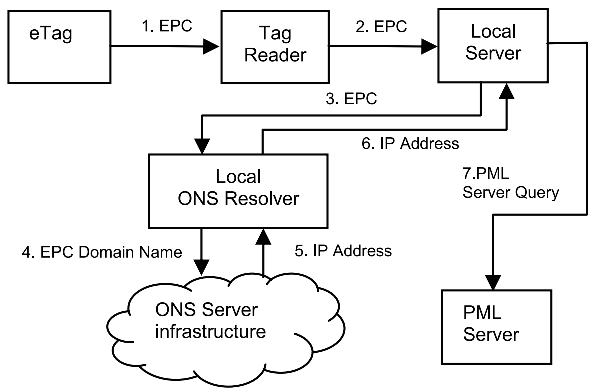

Fig. 3: A tag reader detects the EPC stored in the electronic tag and retrieves the network IP address of the associated PML server through a query of the ONS resolver, which then directs the request to the authoritative ONS server.

The ONS resolver submits queries to ONS servers to obtain the network location of the PML servers. The query results are "cached" to quickly respond to future queries in the same EPC range.

Figure 3 shows the typical information flow and state transitions involved in a typical ONS query. First a tag reader senses and decodes the EPC stored in an electronic tag. Second the EPC is transmitted to a local server from the tag reader. Here it may be time stamped and cached for retrieval or used directly to acquire object information. Third, in the latter case, the local server needs to acquire the network Internet Protocol (IP) address of the associated PML server. The local server performs the lookup using ONS as follows:

- **ONS Resolver** —The local server queries the local ONS Resolver, which may be a process on the same machine.
- **ONS Server**—The local ONS Resolver sends the ONS query to any number of available ONS Servers.
- **IP Address**—The ONS servers respond with the location of the PML Server, or the location of another ONS server that will lead to a server with authoritative information about the desired PML server.

The ONS is based on the standard Domain Name Service (DNS). Within the ONS System, the EPC is translated into a valid DNS domain name. This domain name is used to obtain the set of IP addresses in the standard DNS manner. Unlike DNS, the ONS has multiple roots corresponding to a public and multiple private ONS hierarchies. The private ONS hierarchies are required to locate local information and services stored for a particular EPC.

7. Physical Markup Language

The Electronic Product Code serves as a reference to information on the computer network. There are, of course, many methods for storing information "on-line." These include properitary and commercial databases, and relational databases, such as Structured Query Language (SQL) file systems. Web pages written in the HyperText Markup Language (HTML) are now one of the most common means of storing digital information. New approaches, such as the eXtensible Markup Language (XML), promise a universal means for structured information.

In order to describe physical objects, we propose a new language specific for that purpose — the Physical Markup Language (PML) [Brock-2001b]. Rather than a new syntax, the PML is based on XML and includes a set of schema describing common aspects of physical objects. Industry specific representations may be “plugged into” the common framework or derived from the shared data.

The Physical Markup Language (PML) is intended to be a general, standard means for describing physical objects with particular emphasis on practical applications, such as inventory tracking, automatic transaction, supply chain management, machine control and object-to-object communication.

Our approach has been to develop a series of core components that span a breadth of applications, independent of a particular industry. We have attempted to characterize common aspects of physical objects. For example, objects are composed of matter. They have physical properties, for example volume and mass. They exist in certain forms (solid, liquid and gas) and have measurable states (temperature, pressure and weight). Man-made objects often exhibit structure, having a hierarchical form of system, assembly, subassembly and components. They exist in time and space and are moved about through various means among one or many managers. Products have attributed values of worth, cost, price and ownership.

Although there are many methods to represent common data, for example, time and date, PML strives toward single representations for any data element. The idea is to encourage data translation when encoding or viewing, not during data interchange.

8. Field Trial

The Auto-ID system, including electronic tags, Savants, Electronic Product Codes, Object Name Service, and Physical Markup Language, as well as all of the application software is currently being tested in a large, realworld experiment called the Auto-ID Center Field Trial. The Field Trial is a multi-site, multi-company exercise designed o test and evaluate the fundamental components of the Auto-ID architecture within a supply chain management application.

Test collaborators include the Auto-ID Center, CHEP, International Paper, Gillette, Johnson & Johnson, Kraft Foods, Procter & Gamble, Unilever, Wal-Mart, Sun Microsystems, Coca-Cola and others. The locations of the test sites are proprietary to those participating in the test. The retail outlets include a Sam’s Club and a Wal-Mart store. (At this time, no consumer homes will be included in the test. If they are included later, they will be added on a limited, volunteer-only basis.)

The Field Trial began 1 October 2001 and is intended to end approximately mid 2003. The Field Trial will transition through four broad phases. The first phase involved tagging and tracking logistics units, i.e., pallets. Radio frequency EPC tags were affixed to pallets, and their movements were recorded at entrances and exits. During the first phase, 97% of all tagged items, tagged pallets loaded with un-tagged product, were identified.

The second phase began 1 February 2002 and includes the tagging of pallets and cases. Figure 4 illustrates the basic supply chain used within the first two phases of the Field Trial. Case level identification increases the number of tagged objects by roughly an order of magnitude, but increases the utility and benefits of the system by at least the same margin. To this point during the second phase, 100% of all tagged items, tagged pallets loaded with tagged cases, have been identified as they move through the system.

The third phase begins on 1 January 2003 and includes the tagging of individual salable items in addition to pallets and cases. Item level identification is the vision and goal of the Auto-ID Center research. The third phase will prove the viability of the

Networked Physical World system. The fourth phase includes performing additional testing in European locations. The results of these phases will be used by the participating companies to generate detailed business plans on the execution and deployment of this technology.

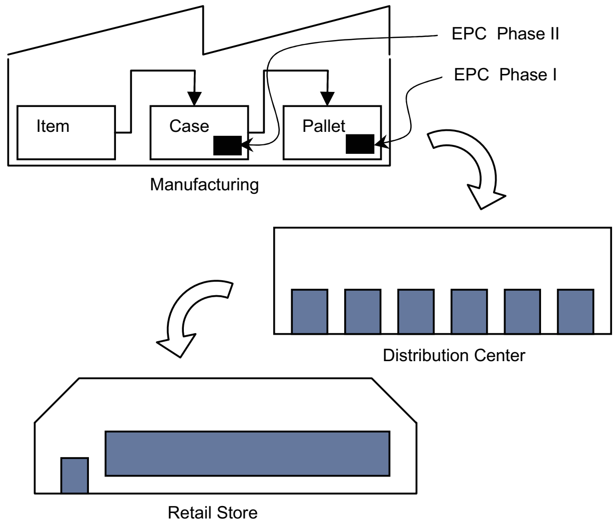

Figure 4: Phase I will track pallets of merchandise provided by the manufacturers participating in the test. The primary objective of this phase is to test and evaluate the ONS and PML software, which were developed by Auto-ID Center researchers working at the Massachusetts Institute of Technology. Phase II will track cases of merchandise. The primary objective of this phase is to test the system's capacity for handling greater volumes of inventory and the information associated with that inventory. Phase III will track products in a retail environment.

The true utility of the Field Trial is to prove the viability of the Networked Physical World system in a realworld (outside the laboratory) environment. The Networked Physical World system has the potential to significantly improve supply-chain management.

Benefits for both commercial enterprises and consumers include better value and better service. For instance, this system will help prevent out-of-stock merchandise and reduce obsolete or out-of-date products in distribution centers. In the end, we believe this technology, fully implemented, will help save billions of dollars of cost within the supply chain.

9. Summary

This paper presents a brief overview of on-going research and development at the Auto-ID Center in 21st century supply chain technology. The Center is a multi-national organization with researchers and labs in four continents. Sponsored by more than 60 companies, government organizations, and trade bodies, the Center is rapidly developing standards, languages, and protocols, as well as experiments and commercial demonstrations, that prove the power and utility of the "networked physical world." The Field Trial is proving the viability of the system components and will prove the economic utility and feasibility of creating "smart objects" that are always connected to the global Internet.

AUTO-ID LABS

References

Auto-ID Center and OATSystems, Sarma, S., Brock, D., Engels, D., Waldrop, J., Putta, L., Ramachandran, S., Nassar, G., The Savant: Version 0.1 (alpha), Technical Report MIT-AUTOID-TM-003, Auto-ID Center, February 2002. http://www.autoidcenter.org/research/MIT-AUTOID-TM-003.pdf.

Auto-ID Center and OATSystems, Sarma, S., Brock, D., Foley, J., Putta, L., Ramachandran, S., Nassar, G., The Object Name Service: Version 0.5 (beta), Technical Report MIT-AUTOID-TM-004, Auto-ID Center, February 2002. http://www.autoidcenter.org/research/MIT-AUTOID-TM-004.pdf.

Barton, J., and Kindberg, T., The Challenges and Opportunities of Integrating the Physical World and Networked Systems, HPL Technical Report HPL-2001-18, 2001.

Brock, D. L., The Electronic Product Code – A Naming Scheme for Physical Objects, Technical Report MIT-AUTOID-WH-002, Auto-ID Center, January 2001. http://www.autoidcenter.org/research/MIT-AUTOID-WH-002.pdf.

Brock, D. L., The Physical Markup Language – A Universal Language for Physical Objects, Technical Report MITAUTOID-WH-003, Auto-ID Center, January 2001. http://www.autoidcenter.org/research/MIT-AUTOID-WH-003.pdf.

Brock, D. L., The Compact Electronic Product Code – A 64-bit Representation of the Electronic Product Code, Technical Report MIT-AUTOID-WH-008, Auto-ID Center, November 2001. http://www.autoidcenter.org/research/MIT-AUTOID-WH-008.pdf.

Cole, P.H., Hall, D.M., Loukine, M, and Werner, C.D. Fundamental Constraints on rf Tagging Systems, Third Annual Wireless Symposium and Exhibition, Santa Clara, February 1995, pp 294-303.

Cole, P.H., and Hall, D.M., Integral backscattering transponders for low cost rf id applications, Fourth Annual Wireless Symposium and Exhibition, Santa Clara, February 1996, pp 328-336.

Esler, M., Hightower, J., Anderson, T., and Borriello, G., Next Century Challenges: Data-Centric Networking for Invisible Computing: The Portolano Project at the University of Washington, Mobicom99, 1999.

Finkenzeller, K., 1999. RFID Handbook: Radio Frequency Identification Fundamentals and Applications, John Wiley & Sons.

MIT Laboratory for Computer Science and MIT Artificial Intelligence Laboratory, Human-Centered Computing: MIT Project Oxygen, Brochure, 2000. http://oxygen.lcs.mit.edu/publications/Oxygen.pdf.

Sarma, S. E., Towards the Five Cent Tag, Technical Report MIT-AUTOID-WH-006, Auto-ID Center, November 2001. http://www.autoidcenter.org/research/MIT-AUTOID-WH-006.pdf.

Uniform Code Council (UCC), Uniform Product Code (U.P.C.) Symbol Specification Manual, January 1986. http://www.uc-council.org/reflib/01302/d36-t.htm.

Uniform Code Council (UCC), Reduced Space Symbology.
http://www.uc-council.org/rss14/.

Uniform Code Council (UCC), Application Standard for Shipping Container Codes (ANSI/UCC6).
http://www.uc-council.org/reflib/00602/d33-t.htm.

Uniform Code Council (UCC), UCC/EAN-128 Application Identifier Standard (ANSI/UCC4).
http://www.uc-council.org/reflib/00403/d30-t.htm.