Published October 1, 2000. Distribution restricted to Sponsors until January 1, 2001.

WHITE PAPER

The Networked Physical World

Proposals for Engineering the Next Generation of Computing, Commerce & Automatic-Identification

Sanjay Sarma, David L. Brock & Kevin Ashton

MIT AUTO-ID CENTER MASSACHUSETTS INSTITUTE OF TECHNOLOGY, 77 MASSACHUSETTS AVENUE, BUILDING 3-449G, CAMBRIDGE, MA 02139-4307

ABSTRACT

The Auto-ID Center at the Massachusetts Institute of Technology is a new industry sponsored lab charged with researching and developing automated identification technologies and applications. The Center is creating the infrastructure, recommending the standards, and identifying the automated identification applications for a networked physical world. All technologies and intellectual property developed at the Auto-ID Center are freely distributed. This white paper outlines the Auto-ID Center's key conclusions and research progress after its first year of research.

Published October 1, 2000. Distribution restricted to Sponsors until January 1, 2001.

WHITE PAPER

The Networked Physical World

Proposals for Engineering the Next Generation of Computing, Commerce & Automatic-Identification

Biography

by Sanjay Sarma
Research Director

Sanjay Sarma received his Bachelors from the Indian Institute of Technology, his Masters from Carnegie Mellon University and his PhD from the University of California at Berkeley. In between degrees he worked at Schlumberger Oilfield Services in Aberdeen, UK, and at the Lawrence Berkeley Laboratories in Berkeley, California. Dr. Sarma's Masters thesis was in the area of operations research and his PhD was in the area of manufacturing automation. From 1995 to 1999, Dr. Sarma was an assistant professor in the Department of Mechanical Engineering at the Massachusetts Institute of Technology. He is now an associate professor.

by David L. Brock
Co-Director

Dr. David Brock received Bachelors degrees in theoretical mathematics and mechanical engineering from MIT, and his Masters and Ph.D. degrees from the Department of Mechanical Engineering at MIT with an affiliation to the Artificial Intelligence Lab. He is currently a visiting scientist at the Distributed Intelligent Systems center at MIT. Dr. Brock is also Founder and Director of Brock Rogers Surgical, a manufacturer of micro- robotic devices. Dr. Brock has worked with a number of organizations including the Artificial Intelligence Laboratory, the Massachusetts Eye and Ear Infirmary, DARPA, Celadon, Loral, BBN and Draper Laboratories.

by Kevin Ashton
Executive Director

Kevin Ashton is Executive Director of the Massachusetts Institute of Technology Auto-ID Center in Cambridge, Massachusetts, and the Cambridge University Auto-ID Centre in Cambridge, England. He is on loan to MIT from the Proctor & Gamble Company, where he is an Associate Director. He is graduate of University College London and is a Visiting Engineer in MIT's Department of Mechanical Engineering.

Published October 1, 2000. Distribution restricted to Sponsors until January 1, 2001.

WHITE PAPER

The Networked Physical World

Proposals for Engineering the Next Generation of Computing, Commerce & Automatic-Identification

Contents

1. INTRODUCTION

The Auto-ID Center at the Massachusetts Institute of Technology is a new industry sponsored lab charged with researching and developing automated identification technologies and applications. The Center is creating the infrastructure, recommending the standards, and identifying the automated identification applications for a networked physical world. All technologies and intellectual property developed at the Auto-ID Center are freely distributed. This white paper outlines the Auto-ID Center's key conclusions and research progress after its first year of research.

2. OUR VISION

The Auto-ID Center envisions a world in which all electronic devices are networked and every object, whether it is physical or electronic, is electronically tagged with information pertinent to that object. We envision the use of physical tags that allow remote, contactless interrogation of their contents; thus, enabling all physical objects to act as nodes in a networked physical world. The realization of our vision will yield a wide range of benefits in diverse areas including supply chain management and inventory control, product tracking and location identification, and human-computer and human-object interfaces. Our vision of ubiquitous automated identification technologies and their applications drives our research agenda and goals.

3. KEY CONCLUSIONS

The scale of our vision impacts all aspects of our research and weighs heavily upon our key conclusions after our first year of research. A networked physical world requires a system that includes all traded objects in the world. Such a system must scale to unprecedented proportions, quickly becoming the largest man-made system ever. For perspective, we expect such a system to ultimately handle more than one trillion (10^{12}) new unique objects annually [1].

[1] Today, in excess of 5 billion bar-coded items are scanned every day. This amounts to 1.7 × 1012 bar-coded items annually.

3.1. A single, open architecture system for networking physical objects is more valuable than multiple, smaller scale, alternatives.

The question of scale may be avoided by designing isolated, independent systems for different classes of objects – for example, one system for books and another for food. However, this fragmentation greatly limits the value of the activity, increases its overall complexity, increases costs, and makes successful proliferation far less likely.

A well known parallel to our networked physical world vision is the Internet. The Internet is successful because it is freely available to everybody and readily scales with the addition of new nodes. This has made networking ubiquitous and, therefore, valuable. Private networks, such as the original America Online, or Apple's "Apple World", that did not offer the same openness were quickly displaced or subsumed by the Internet. Metcalfe's Law explains that the value of a network is the square of the number of users it has. When we apply the law to a network of physical objects of the scale required in our vision, it is clear that a single, open architecture for networking physical objects is much more valuable than smaller scale alternatives.

Designing this system raises a number of interesting research challenges and requires careful engineering.

3.2. The open architecture system must be platform independent and highly interoperable.

"Objects" is a broad term covering all possible combinations of size, shape, material, and implementation. No single embedded tag technology can perform adequately with all types of object. In addition, government regulations vary by region, prohibiting different contactless technologies in different parts of the globe. Thus, the open architecture system must be platform independent and highly inter-operable. Again, this is supported by the success of the Internet, which can function with any type of compatible hardware: PCs, Macintoshes, UNIX platforms, and now other devices such as cell phones, televisions, cable set-top boxes and even washing-machines.

3.3. The open architecture system must require a minimum of performance from the tag technology embedded in the objects.

Many objects, for example, food items, are consumable and very sensitive to additional costs. These cost sensitive objects must be included in all open architecture systems that implement our vision since we do not exclude any object, or class of objects, from our networked physical world vision. By physically tagging every object, cost savings may be obtained through the better management of the supply chain. These savings are such that the cost of the physical tag must be a few pennies or less if it is to be used on a cost sensitive object. The maximum additional cost acceptable to cost sensitive objects requires the use of electronic tag technologies that (a) meet the maximum cost constraint and (b) meet the open architecture system performance constraints. Since higher technological performance is typically associated with higher cost for all electronic technologies, constraints (a) and (b) lower the minimum performance of the tag technology that should work with the open architecture system.

3.4. The open architecture system must be flexible and adaptable to change and improvement long after it has been installed.

An installed open architecture system is, by definition, a legacy system. A legacy system deployed on the scale we envision will be costly and unlikely to be replaced. Further, any technology deployed today must be assumed to become obsolete in a few years. Consequently, it is imperative that the initially installed system be amenable to new technologies and system upgrades without requiring the entire system to be replaced.

4. SYSTEM OVERVIEW

We have developed an open architecture system that enables the implementation of our vision of a networked physical world. Our system is based upon our key conclusions. In this section, we describe the primary components for our system.

4.1. The Tag Data: The EPC Code

The critical design feature for our open architecture system is the separation of the information about an object from the object. Maintaining information separate from an object enables minimal performance tags to be used while increasing the robustness, scalability, and flexibility of the system.

After making this design decision, the question is then 'How can the data stored on the tag be used to locate the object information?'

We propose assigning a wholly unique, searchable idnetification number to each physical object. We call this number an EPC, or **Electronic Product Code**. This number is analogous to the IP (Internet Protocol) address given to compute nodes on the Internet. It is also somewhat similar to the UPC/EAN (universal product code/international article number) system, although where UPC/EAN identifies types of object, the EPC uniquely identifies each individual object.

Figure 1: Electronic Product Code

ELECTRONIC PRODUCT CODE

X . XXX . XXX . XXXXX

Header — Manufacturer — Product — Serial Number

Given our desire for a unique searchable number for every object, we propose a simple scheme for the electronic Product Code: a 96-bit scheme with an 8-bit header and three data partitions, as shown in the above figure. Each 'X' in the figure indicates 8 bits. A 96-bit number was chosen as a compromise between the desire to ensure that all objects have a unique EPC and the current limitations of our proposed tag technology described below. The partitions of the EPC enable a hierarchical search for the information about the object tagged with a particular EPC.

The EPC header is used to indicate the format, or total length and field partitions, of the EPC. The header is essential for the flexibility of the system. It permits multiple EPC formats; thereby allowing the seamless use of longer bit-length tags as the technology matures. The header also permits the bits of a particular length to be redistributed. For example, a longer manufacturer number (with a corresponding shorter length product and serial number fields) may be used for producers that have few products. This follows the model of IP address classes.

Following the tradition of the original Uniform Product Code, we assign the next two data partitions to manufacturer and product code. Although the UCC-12 code provides up to 100,000 manufacturers and products, we propose an increase to approximately 268 million manufacturers, using 28 bits, and 16 million product code assignments, using 24 bits.

The final data partition represents the product serial number, allowing each tagged item to be uniquely identified. A 36-bit serial number allows over 68 billion uniquely identified objects for each of the 16 million products. Taken together, each manufacturer is allowed 1.15×10^{19} uniquely identified items, which should be sufficient for the indefinite future.

4.2. The Physical Tag

The EPC code comprises the required data to be stored in a tag. We are proposing that the physical tag itself be an Electromagnetic Identification (EMID) tag. An EMID tag is a memory device with 'circuitry' for communicating wirelessly with an external tag reader. Tags can be classified by technology and by performance in many ways. The important characteristics are transmission technology, modulation scheme, encryption scheme, fabrication substrate, transmission frequency, anti-collision algorithm, read range, tag size, number of bits, individual addressability, read-write ability, and power source. Appendix A provides a summary of these considerations.

The most important objective in the design of tags for our application is cost. To keep costs down, it is necessary to design the overall system in such a way that the requirements we place on the tags themselves are minimal. Reducing tag memory by minimizing required data reduces tag cost. Similarly, reducing required tag functionality, such as read/write capabilities, reduces tag cost. It is possible to write on some EMID tags today, but writing requires extra features and memory that make tags more expensive. In order to meet our preliminary cost target of 10 cents, we circumvent the need for writing by requiring only that the EPC be stored on the tag – information pertinent to a tag can then be stored on a networked database resource. In the information age, data need not travel piggy-back on the product – it can travel separately on the Internet.

4.3. The Key Tag Characteristics

There are three tag characteristics that require special attention. The first characteristic is the frequency. There has been some question as to whether a single, standard frequency should be chosen. Four carrier frequencies receiving early attention as representative of the low, high, and ultra high frequencies are 125 kHz, 13.56 MHz, 900 MHz and 2.45 GHz. Unfortunately, each frequency has a problem that prevents it from being universally applicable. In general, lower frequencies have better penetration of materials, but higher frequencies have better data-rates. With line of sight, high frequencies also have greater ranges – but they require more power. Additional considerations in the selection of a frequency are compliance with various national and international regulations. The situation is further complicated by the constraints placed on transmission power in different countries. Although most systems are likely to fall under the levels where power regimes are regulated, long range systems may require higher power. Thus, the frequency issue is also tied to application. Finally, the reasons for selecting these frequencies are political, not technical: they are generally designated as ISM bands around the world; therefore, they can be used by tags without regulatory impediment in one or more regions.

It is our view that it is premature to answer the frequency question at this point. There are three reasons for this decision: 1. At the scale we are proposing for the deployment of tags, it is likely that tags will find new and unexpected uses, the requirements for which we cannot know now. 2. There is promise that it will be possible to design readers that can read a wide range of frequencies, from the kHz to the GHz ranges. In fact, one of the projects for 'wide-band' or agile readers is being led by Professor Neil Gershenfeld of MIT, an associate of the Auto-ID Center. 3. And, 3. many of the advanced communication technologies, such as spread spectrum communication, show promise of being implemented in EMID tags in the near future.

The second key tag characteristic is the modulation scheme. There are at least a dozen classes and sub-classes of modulation schemes. The most important categories of modulation schemes are Amplitude Shift Keying (ASK), Frequency Shift Keying (FSK), Phase Shift Keying (PSK), and multiple access schemes (TDMA, FDMA, CDMA and WCDMA). We are researching these issues. It is not clear which is a preferable standard at this point.

The third key tag characteristic is anti-collision, or the ability to read many tags at once. Professor Sunny Siu has developed a 'hands up' protocol that achieves close to the theoretical maximum performance. This protocol is also memory-less – and requires little additional functionality in the tag. We are currently working with tag manufacturers to see if this protocol can be implemented cheaply. If we are successful, it may form the basis of a future standard recommendation.

In summary, we recommend no action on a frequency standard at this time. We are studying modulation scheme and will recommend a course of action soon. We have developed an anti-collision scheme that is currently being considered for implementation and testing.

4.4. Local Networking Technology

A world with trillions of tags will require millions of readers – in warehouses, loading docks, factory floors, trucks and shelves. This is a very different world than one in which a few hand-held units are used to scan items manually. Large numbers of readers, and additional sensors, will need to be networked. We are developing inexpensive ways of networking readers, sensors, and local databases.

Figure 1: Tags and Readers connected via local database

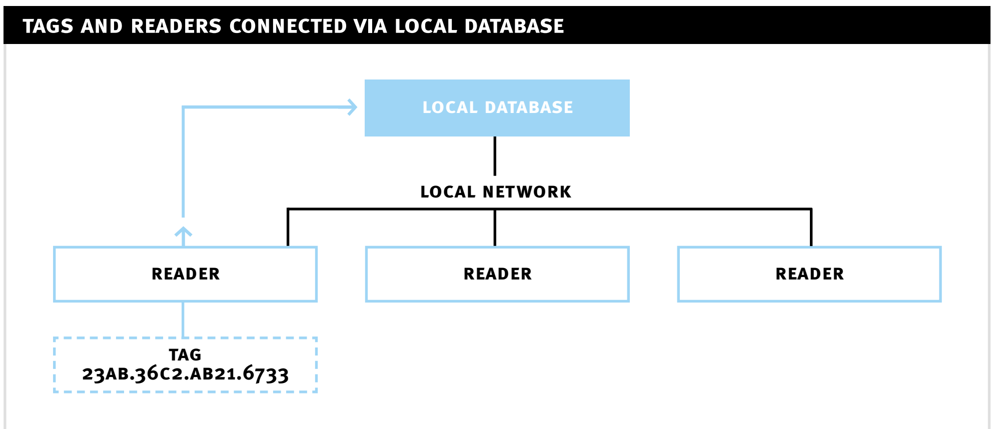

The most attractive way to network readers and sensors is through a TCP/IP network. However, TCP/IP controllers are expensive today. We are looking at other standards including RS 232 (implemented), RS 485 (implemented), Lon Works IEEE 1473 (under investigation). Issues being investigated include data rate, power requirements, wire-length considerations, and robustness.

The figure above shows that the tag readers will be connected to a local database. Information will be stored in the database at a regular 'tick rate' that will depend on the application. We are developing inexpensive implementations for the database. Our current implementation uses public database software called MySQL and is implemented on a low-end PC with RAID storage. The database can be accessed from the Internet by an entity with the proper authorization.

Several emerging standards also bear watching. BlueTooth™ is an emerging RF LAN standard championed by Nokia, Ericsson, IBM, Intel, Microsoft, and Toshiba. We are investigating this standard, as well as acquiring the application kits for some initial testing. Jini™ is a discovery and registration protocol developed by Sun Microsystems, one of the sponsors of the Auto-ID Center. We have already implemented and tested Jini™ at the Center. As Jini™ evolves, we will work with Sun to ensure compatibility. Jini™ is an example of a plug and play protocol. Microsoft is also developing a protocol called Universal Plug n' Play (UPnP). We are involved in the UPnP working groups and intend to work towards understanding this standard as well. A summary of these standards is presented in Appendix B. Standard systems to integrate reader networks and databases locally will facilitate faster adoption.

4.5. The Object Name Service (ONS)

Since the EPC is the only required data stored on a tag, it must be used to find additional data about an object. This additional data may be stored on a server connected to the local network or to the Internet. The server is identified via a free look-up system which we call ONS, or **Object Name Service**. This is analogous to the DNS, or **Domain Name Service**, used to locate information resources on the Internet.

The first generation prototype of the ONS was completed last year. Since then, we have extended the design of the ONS to make it more scaleable and to provide greater flexibility to users. The second generation ONS specification and implementation are currently underway. Along with the ONS, the specification includes descriptions of the functionality of resolving and caching at the client side. We are designing the system to handle traffic on the order of 10^8 transactions per day. ONS is the 'directory service' for connecting the physical object and the information about that object.

4.6. The Physical Markup Language (PML)

The information about an object will be specified using the Physical Markup Language (PML), a version of the popular XML meta-data language. XML is a natural choice for our system because we require a human-readable test format that is also a database standard supported by the world-wide web. The further advantage is that queries can now be stated in XQL format. XQL is an emerging XML database querying standard now under consideration by the W3C. XQL is similar to the relational database standard SQL, and the output of XQL is also an XML file.

The syntax and semantics of PML must be administered by a user-oriented body such as the UCC in conjunction with the user community. Currently, we are developing some technical guidelines for the PML standard. Our first definition was a simple example for food items, and consisted of FDA nutritional information guidelines. The next example will be in the area of logistics. The UCC is currently engaged in modeling business processes using UML. Eventually, the job of defining PML semantics should dovetail with the UCC's modeling group. The lingua franca for describing objects in the system will be PML.

4.7. The PML Servers

PML files are stored in special servers called PML servers. The manufacturer of an object will maintain a PML server for all of its objects. Currently, our system uses HTTP (Hyper-Text Transfer Protocol). We are now considering a simpler protocol and are evaluating an UDP based system. Because of where the Internet will be in two years, we feel that a lighter protocol can be used effectively.

The PML server will store a variety of different types of information: class data (such as the name of the product); instance data (such as the date of manufacture of a particular item); distributed data (such as where an item is located); and machine readable instructions and algorithms (such as how to cook the item if you are a microwave oven). The reason for the classification of information is because the implementations of each of these types of databases may well be different for efficiency and compression. The class information will be common and can be stored on standard databases. Instance information is specific, but also very copious. A pre-processed, custom implementation may be necessary.

Whatever the implementation, the PML server must serve text PML files. Therefore, the XML/XQL queries must be translated to the specific database format. In this sense, the XML/XQL handler can be seen as middleware to the actual implementation of the database. The PML server stores object specific information, which it will serve as PML files upon an XQL request.

4.8. Application Software

In the coming years, hundreds of applications will evolve for automated identification using contactless tag technologies. Like the Internet, the infrastructure we have described here is a platform for enabling these tasks. Some core functionalities are of immediate interest, however. They are:

1. Tracking
2. Locating
3. Pattern matching and rule-checking
4. Inference
5. Modeling

We are developing a protocol called RSVP for tracking items sequentially as they move from stage to stage in a supply chain. Brute-force tracking searches must be avoided at all costs because the volume of transactions this will spawn will overwhelm the network. Generalized searching will be implemented using a publish and subscribe architecture instead of active searching. Pattern matching, rule checking, and inference software will operate as daemons on the networks. With the appropriate access, these daemons will be able to check for certain conditions – such as an item being stocked in the wrong shelf. At the same time, modeling functions can act as watchdogs that evaluate whether or not the circumstances in a system are evolving as they should. In essence, enabling these functionalities is the primary goal of the system we have described.

5. RESEARCH CHALLENGES

We now summarize the challenges before us. We believe that all of these challenges can be met from both a price and performance point of view. The important questions relate to time, cooperation, and organizational commitment. We list below the more pressing issues.

Tag cost:
Tags with a range of a meter and a cost of 10 cents are not commercially available today. We do, however, believe that tags at this cost can be achieved with concerted engineering effort. The difficulties in manufacturing tags are related to the silicon cost, the packaging costs, and antenna costs. At MIT, we have the technology to address each of these areas with innovative, aggressive solutions. We believe that with an initial outlay of a few million dollars, it will be possible to invest in the optics and handling technology necessary to bring down the silicon costs. Handling costs too can be controlled with aggressive precision machine design. Finally, antenna costs can be addressed with new manufacturing processes such as lithography material printing.

Tag/reader performance:
The predictability and robustness of wireless communications using low cost tags needs improvement. We believe this can be addressed with improved Integrated Circuit design, refinements in tag packaging, and developments in reader technology. In addition, it is possible to incorporate error correction and redundant, multiple read opportunities into the system and its software.

The frequency question:
A Masters thesis currently under way in the Media Lab at MIT will describe a reader that can read multiple frequency tags. This a matter of good digital design. The only technical challenge appears to be that of a wide-band antenna. However, several groups are working on promising approaches to wide-band antennas.

6. NEXT STEPS

In preparation for field testing, we are proceeding on four fronts.

System development:
Having prototyped the important elements of the architecture, we are now entering the development phase. ONS, PML, the PML server, and the local system are now being implemented. We are also writing Internet Engineering Task Force (IETF) papers to transition these systems into Internet standards.

Tag technology investigation:
We are working on several aspects of tag technology. First, we are working on manufacturing issues for tags to bring unit costs down. Second, we are working on modulation schemes to evaluate potential standards. Third, we are monitoring developments on the frequency question. Fourth, we have developed a new anti-collision algorithm and are looking at possible implementations. Finally, we are studying the electro-magnetics of tags, readers, and radars.

Applications:
The goal of this system is to open up new applications. We currently looking at three categories of applications. First, we are looking at warehousing and logistics applications. A recent Master's thesis studied the theoretical impact of tagging on the bull-whip effect. Second, we are looking at industrial applications. For example, we are researching the use of tags in industrial robotics and picking systems. Finally, we are looking at "utility" systems – ranging from a glove that identifies an object being picked up to making inventory information available over WAP-enabled mobile devices.

Incorporation of legacy systems and existing related activities:
The generality and flexibility of the system we propose is so great that existing legacy systems and prior activities could be incorporated as subsets of the general system with relatively little modification. Current numbering systems, for example, can be handled as part of an item's PML file, as can labeling requirements. We propose continuing activity to ensure that key legacy systems are identified and, where necessary, harmonized or provided for.

7. APPENDIX A: BRIEF OVERVIEW OF TAGS

Tags are devices with non-volatile memory that can be interrogated remotely. Usually, radio frequency (RF) tags have memories ranging from a few bits to several kilobytes and communicate with the tag reader at frequencies ranging from 100 kHz to several GHz. Tags can be classified in several ways, such as by the transmission technology, the transmission frequency, the range, whether the tag is powered or not, the size of the memory, read-write capability, and anti-collision algorithm. We examine these criteria below. We are aware of more than fifty RF tag manufacturers world-wide, and there are several combinations of technologies available.

7.1. Transmission Technology

There are three forms of energy, or signal transfer, that may take place between readers and tags. The first is power transfer from the reader to the tag for the operation of the tag. The second is data sent from the reader to the tag. The third is data sent from the tag to the reader. Power and information transfer may occur through either a near field coupling or far field energy harvest. There are several types of near fields that can be generated for coupling. For example, Arizona MicroChip's MicroID™ tags couple to a magnetic field via inductive coupling, while Motorola's Bistatix™ tags couple to the electrostatic field via capacitive coupling. Loosely stated, the near field types are capacitive and inductive. Each phenomenon has its limitations in terms of range and frequency.

7.2. Frequency

There are three ranges of frequency in which RF tags tend to operate. They are: low frequency, 30 – 300 kHz; high frequency, loosely used to refer to the range 300 kHz – 300 MHz; and ultra high frequency, 300 MHz – 3 GHz. There are several factors that affect decisions on frequency. The first is range. Low frequencies typically require near field coupling and have smaller ranges, but also are absorbed less by most non-conductive materials and by water. Ultra high frequencies typically require far field energy harvesting and have higher range and offer better bandwidth, but suffer the disadvantage that because they are less likely to pass through materials, and because they tend to be more "directed," they require line of sight. Furthermore, ultra high frequency transmission needs more power. Finally, there is the issue of regulation. In the US, frequency questions are answered as much by FCC regulations as by technological questions. A map of the US frequency map is shown on the following page.

7.3. The Question of Power

Tags can be active, passive, or semi-active. Active tags have a battery for two reasons: for keeping the tag circuitry running and for communication with the tag reader. Passive tags have no battery. Instead, they obtain power from the reader or from an ambient field. Semi-active tags have a battery on board to keep the tag circuitry running, but use reader power for communication. Power impacts cost and, to some extent, robustness of the tagging technology. Powered tags have a limited life and cost far more than passive tags.

7.4. Read-Write Capability

The cheapest tags are read-only. The tag number is usually stored in ROM or is hardwired in the circuit by cutting connections using a laser. There are more flexible versions of ROM, one of which is the EEPROM. EEPROM modules can be embedded inside tags and can be over-written tens of thousands of times. However, EEPROMs are expensive and require more power than simple hardwired circuits. This increases tag cost and reduces range. Newer technologies, such as FPROM's, offer hope for low power read-write capabilities because they require less power and operate faster than EEPROMs. However, manufacturing problems have curtailed the user of FPROM's thus far. SRAM's can also be used on tags but they require uninterrupted power. Therefore, tags that use SRAM's require batteries.

7.5. Anti-Collision

When multiple tags respond simultaneously to a reader, it is called a collision. There is a need to disambiguate these collisions, and this procedure is known as anti-collision. There is an interesting trade-off between cost, frequency, and range. Most anti-collision algorithms require significant computation on the microchip of the tag. This increases the size of the chip, increasing cost and, because of higher power consumption, reducing range. The higher frequency systems can resolve more collisions per second. However, higher frequency systems also dissipate more energy. The most efficient anti-collision algorithm would be one in which each tag is assumed to be able to hear every other tag in the same reader's vicinity. However, this too reduces the effective range of the reader-tag system. Finally, most anti-collision algorithms require some amount of memory on-board that retains information for the duration of the collision resolution process. Such memory requirements can also increase cost and power requirements. We have developed a memory-less anti-collision system that requires little additional computing on the tag while ensuring near maximum (theoretical) performance.

In summary, there are literally hundreds of choices in the design of tag-reader systems. Tag technologies are very capable today; tags available on the market today can achieve tens of meters of range, and several kilobytes of memory. It is tempting to gravitate towards a single, high-technology tag which can 'do everything.' However, the truth is that simple, read-only tags remain the most promising technology that could one day become cheap enough to proliferate. A tagging system and standard that can accommodate the entire range of tags, from read-only tags to high-end tags, has the greatest chance for success.

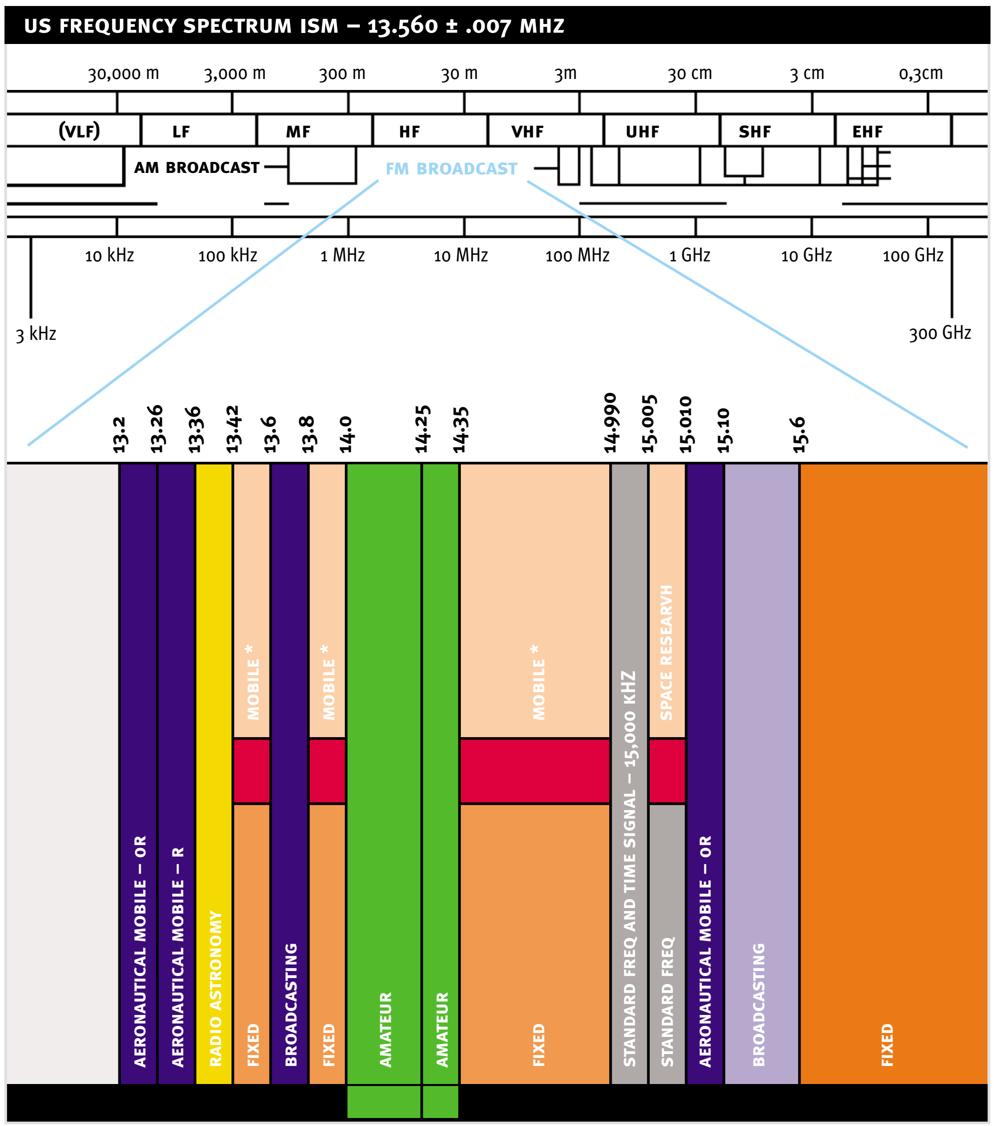

Figure A: A zoomed in picture of the US frequency spectrum in the vicinity of the 13.56 ISM (Industrial-Scientific-Medical) band.

8. APPENDIX B: AN OVERVIEW OF LOCAL NETWORKING SYSTEMS

Several technologies and standardization efforts in networking and communications have emerged recently which may affect the Auto-ID Center. We summarize the better known standards below, and then describe the common functions of most of these technologies.

8.1. Resource Discovery and Registration: Jini™

Jini™ connection technology is based on a simple motto: that devices should work together without the need to find drivers and without facing problems of incompatibility among operating systems. Jini™ enables computers and devices to form spontaneous networks. In this **federation** of devices – including computers, printers, and sensors – devices need simply to be plugged into the network to become functional. The system works as follows. Each device provides an interface to **services** that may be used by other devices in the community. Devices and their corresponding services are registered with a local **lookup service**. When a new device is plugged in, it first embarks on a **discovery** of the local **lookup service** and **joins** the federation. Jini™ technology also defines a mechanism, which is called the **leasing and transaction** mechanism, to provide resilience for temporary associations in a networked environment.

8.2. Plug and Play: Microsoft's Universal Plug and Play

"Universal Plug and Play (UPnP) is an architecture for pervasive peer-to-peer network connectivity of PCs of all form factors, intelligent appliances, and wireless devices. UPnP is a distributed, open networking architecture that leverages TCP/IP and the Web to enable seamless proximity networking in addition to control and data transfer among networked devices in the home, office, and everywhere in between."
– From the UPnP web-site at www.upnp.org.

8.3. Wireless Lan: Bluetooth™

Bluetooth™ wireless technology operates in a globally available 2.4 GHz band. The Blue-tooth™ specification answers the need for short-range wireless connectivity within three application areas: real-time voice and data transmissions, elimination of the need for numerous and often proprietary cable attachments, and ad-hoc networking. The specification supports both point-to-point and point-to-multi-point connections up to 7 active slave Devices can communicate with an active master Device. While the number of Devices is actually unlimited, only eight Devices can be active at any given moment. Bluetooth™ provides data security by encryption, authentication, frequency hopping, and automatic power adaptation to reduce required range. It provides protection from interference in noisy environments by advanced error-correcting methods. It supports data transmission rates of 1Mbits/sec and a range of approximately 10 meters that can be extended to over 100 meters with the use of amplifiers.

8.4. Other Services

Other emerging standards include the HAVi standard for home multi-media connectivity, SLP (Service Location Provider), a standard under development by the IETF for self configuring service discovery, and Salutation, another similar effort that can be studied at www.salutation.com. A concise summary of these protocols is provided by the IBM Web site at http://www-3.ibm.com /pvc/nethome/

networking.shtml. Most of these protocols have some common elements. The basis for all of them is a client agent and a service agent. The client agent is a software component that runs in a device (for example, a tag reader) and searches the network for services such as a controller and a database. The service agent is for devices that provide services to other agents. A registry is an area where clients and service providers advertise their capabilities and needs. The discovery mechanism is that part of a service discovery protocol that specifies how a client locates the service discovery infrastructure such as a registry. Finally, once a client has located a service and has successfully negotiated access to the service, it must determine (and perhaps acquire) the protocols to actually interact with the service (for example, IPP, LPR, HTTP, FTP, Java RMI).

9. REFERENCES

In addition to research at the Auto-ID Center at MIT, the material in this paper is derived from several sources. We list them below.

1. **Kundapur, Niranjan. On Integrating Physical Objects with the Internet.**
   Master Thesis, Department of Mechanical Engineering, MIT. 2000.

2. **Finkenzeller, Klaus. RFID Handbook; Radio Frequency Identification and Applications.**
   John Wiley & Son, Ltd., West Sussex, England. 1999. Comment: An excellent reference for automatic identification technologies.

3. **IBM website:**
   http://www-3.ibm.com/pvc/nethome/networking.shtml.
   Comment: an excellent summary of today's emerging small-area networking protocols.

4. **UPnP website:**
   http://www.upnp.org/. Comment: The Universal Plug and Play Forum.