SpotON: An Indoor 3D Location Sensing Technology Based on RF Signal Strength

Jeffrey Hightower and Gaetano Borriello
{jeffro, gaetano}@cs.washington.edu
University of Washington
Computer Science and Engineering
Seattle, WA 98195

Roy Want
want@parc.xerox.com
Xerox Palo Alto Research Center
Palo Alto, CA 94304

UW CSE Technical Report #2000-02-02
February 18, 2000

Abstract

Providing a reliable technology and architecture for determining the location of real world objects and people will undoubtedly enable applications, customization, and inference not currently possible. This paper documents the creation of SpotON, a new tagging technology for three dimensional location sensing based on radio signal strength analysis. Although there are many aspects to the topic of location sensing and several of them will be briefly discussed, this paper is primarily concerned with the hardware and embedded system development of such a system.

1 Introduction

In the next century, consumer computing will rely heavily upon mobility, context awareness, task inference, and distributed services. This is not simply a problem of ad-hoc networking of existing applications, but one that requires the development of new devices and services to measure, record, and leverage the state of the physical world. It is also critical that this real world awareness must have a low cognitive load to the user, thus creating a level of convenience in applications that does not exist now. The terms invisible and ubiquitous computing captures these principles.

While there are certainly many facets to invisible computing, location-sensing systems capture our particular interest and are a fundamental technology in the field. The overarching goal of SpotON is to create and analyze a fine-grained indoor location-sensing system and the associated services for use within an invisible computing framework. This paper documents the process of creating the SpotON object tagging technology based on radio signal strength analysis and focuses mostly on the hardware and embedded systems aspects. Such an approach combines the advantages of wireless location systems (fine granularity) with that of infrared-based systems (detection at a distance).

This paper is organized as follows. Section 2 presents relevant background information and a review of related work. Section 3 documents the RFIDeas experiment – our initial attempt to develop a signal strength based location sensing system using commercial components. Section 4 describes our custom hardware design based on the lessons learned from the RFIDeas experiment. Finally, we speculate about applications, discuss future work, and offer conclusions in Sections 5, 6, and 7.

2 Background

2.1 Positioning, Tracking, and Location Sensing

Before a general survey of similar work is presented, it is necessary to offer a definition. We wish to partition systems into two categories: positioning and tracking. Positioning systems are those that provide the means to determine location and leave it up to the user to compute his or her actual position. Tracking systems, on the other hand, are those that monitor objects in their purview without involving the tracked objects in the computation. Many systems are somewhat of a hybrid of these two types – whether by design or simply by configuration – often in the interest of security and privacy. We apply the term "location sensing" to such hybrid systems including our own where the policy of manipulating location data attempts to be separate from the mechanism of actually pinpointing the object.

2.2 Related Work

There have been many systems and architectures over the years tackling the problem of determining location. Since each was developed to fulfill a different goal, they vary widely in many parameters including accuracy, cost, size, configurability, security, and reliability. Examples include GPS, Active Badges and the Xerox ParcTAB, AT&T Cambridge Ultrasonic Bats, Microsoft Research's WaveLAN system, the Smart Floor from Georgia Tech, PinpointCo's radio tags, various computer vision systems, and various cellular phone based solutions.

The Global Positioning System (GPS) provides an excellent architecture for determining location and is perhaps the canonical example of a positioning system under definitions presented in section 2.1. The worldwide satellite constellation provides reliable and ubiquitous capability and, assuming a differential reference, allows transceivers to compute their location to within 1-5 meters.[1] The difficulty with GPS is that it functions minimally or not at all when the receivers are indoors.

The first and arguably archetypal indoor location sensing system is Active Badges developed at AT&T Cambridge.[2] A small infrared beacon is worn by every person and emits a globally unique identifier every 10 seconds. A central sever collects this data from fixed IR sensors around the building, aggregates the data into a central repository, and provides an API for applications to take advantage of the data. An extension to this work used by the Xerox ParcTAB system implemented a 360-degree infrared "deathstar" to address the problem of IR directionality.[3]

Many groups have explored using computer vision technology to implement a tracking system. Easy Living and Perceptual User Interfaces (PUI) are two Microsoft Research groups working in this space. Easy Living attempts to outfit a home environment with stereo vision technology.[4] The PUI group focuses on enhancing desktop applications with extra input provided by inexpensive camera technology often present on consumer PCs.[5] Vision system are perhaps the canonical example of tracking systems using our definition from section 2.1. The objects seen by the cameras do not usually participate in the location computation.

Another group at Microsoft Research has developed a building wide tracking system based on WaveLAN wireless networking technology. Their system leverages the signal strength and signal to noise ratio available from the WaveLAN network interface card (NIC) and triangulates 2D position within a building using either empirical data or a mathematical model of indoor radio propagation. Advantages of this approach are that it requires very few basestations and uses the same infrastructure that provides general wireless networking in the building. The difficulty is that the object being tracked must support a WaveLAN NIC, which may be impractical on small or power constrained devices. In addition, it is not trivial to generalize this approach to multi-floored buildings or three dimensions. The system can place objects to within 3 meters of their actual position with 50% probability.[6]

PinpointCo sells a commercial product called 3D-iD to perform indoor location tracking based on RF code and phase.[7] 3D-iD is similar to both our work and the MSR WaveLAN system. It had decent accuracy and is somewhat scalable but has the disadvantage of being very expensive.

AT&T Cambridge has more recently developed an ultrasonic tracking technology augmenting the previous generation Active Badges to provide more accurate positioning in certain circumstances. Users and objects are tagged with “bats” that emit periodic ultrasonic signals to receivers mounted throughout the ceiling. This system is very efficient and has the added benefit of being able to extract orientation information in addition to basic position data.[8] Unfortunately, using ultrasonics in this way requires a very large fixed infrastructure throughout the ceilings and is rather sensitive to the precise placement of the ceiling sensors thus its scalability is questionable.

The US Federal Communications Commission’s E911 initiatives call for wireless phone providers develop a method to effectively locate phones placing emergency 911 calls.[9] These rulings have spawned many companies investigating technology to determine location using the cellular telephone infrastructure because in addition FCC compliance, cell phone location will enable new valuable services. For example, a user could query his wireless telephone to request the nearest gas station, post office, movie theater, bus, pr other such service. Various RF techniques including angle of arrival to the antenna array, signal attenuation, and time stamping as well as established technologies such as GPS are being employed in this space. These approaches differ from our work in that they are not focused on fine-grained positioning nor do they address the indoor environment in particular.

There is also a large body of existing work in location tracking in support virtual reality and animation motion capture. Technically, many of these systems can provide valuable insight into developing similar systems for ubiquitous computing. For example, it has been shown that CDMA-like radio technology can be used for precise position tracking (on the order of 2mm grain size) for virtual environments.[10] However, three important issues separate location sensing for invisible computing from most of these systems. First, these systems are often quite expensive and thus not readily deployable in the ubiquitous sense. But more important then cost, many of these systems are not designed to be scalable even to a building wide level – they are designed to capture position well in a single room immersive environment. Finally, VR and animation tracking systems usually do not incorporate sensitivity to the privacy of location information nor do they provide a general management architecture enabling many diverse applications to leverage the location data.

Our idea to develop a fine grained tagging technology based on RF signal strength is most similar to the MSR WaveLAN and the Pinpoint systems, however laboratory experiments indicate that we can achieve better resolution and accuracy then the MSR system with a much lower cost then the product from Pinpoint. Furthermore, we believe accuracy and efficiency could be enhanced even further by the addition of sensor fusion techniques such as integrated accelerometers and online building maps.

3 RFIDeas Experiment

SpotON began with a question: Why build custom tag hardware if suitable components already exist in the commercial market? A search led to an Illinois company called RFIDeas and their AIR ID product. AIR ID is an "adjustable long range active ID badge, reader and software solution for desktop computers. Featuring hands-free login, automatic desktop computer locking based on the user's proximity and read/write memory in the AIR ID badge."[11]

Figure 1: RFIDeas badge (left) and basestation (right).

AIR ID boasts a simple RS232 basestation protocol and a signal strength measurement with "multibit" accuracy. Thus, although it is designed for workstation auto-logon, we felt that perhaps the RFIDeas setup could be used in a different way to provide signal strength measurements in our more robust location-sensing architecture.

3.1 AIR ID Analysis

The first step was to perform a simple experiment to determine if it was reasonable to believe that AIR ID could provide the signal strength data required and warranted further investigation. The experiment was setup as follows:

1. In a large indoor space, a single basestation transmitter is placed at a central point in the room connected to a laptop running our measurement and logging program.
2. 12 badges are placed in a concentric circle of a known radius surrounding the basestation.
3. Several signal strength measurements are taken and logged.
4. All badges are moved to slightly increase (by approximately 1 foot) the radius of the circle around the basestation.

5. Steps 3 and 4 are repeated until the majority of the badges are out of range of the basestation.

Figure 2 below summarizes the results. The Badge Average data points are the average of all twelve badge signal strength measurements as distance from the basestation increases and are in abstract units. As expected, the signal strength drop-off roughly obeys an inverse square drop-off as distance from the basestation increases. The variations are most likely the result of signal reflection and interference inherent in an indoor environment. In addition, we found that badges were fairly uniform and consistent in their signal strength – none seemed to report consistently above or below average signal strengths nor did they drift widely over time.

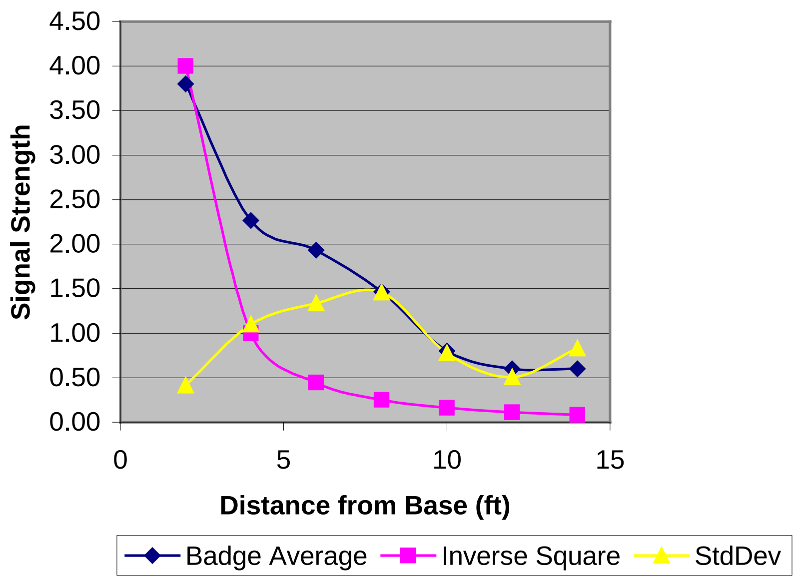

Figure 2: RFIDeas Signal Strength Results

Admittedly, this experiment is very simple, however, these basic results were sufficient evidence for us to believe it might be possible to use AIR ID in the construction of a location sensing system.

3.2 Concept

The location sensing architecture we set out to build is conceptually simple even though the algorithm implementation is certainly nontrivial. Multiple basestations provide signal strength measurements mapping to an approximate distance. A central server then aggregates the values to triangulate the precise position of the tagged object. Finally, the computed object positions are published to client applications.

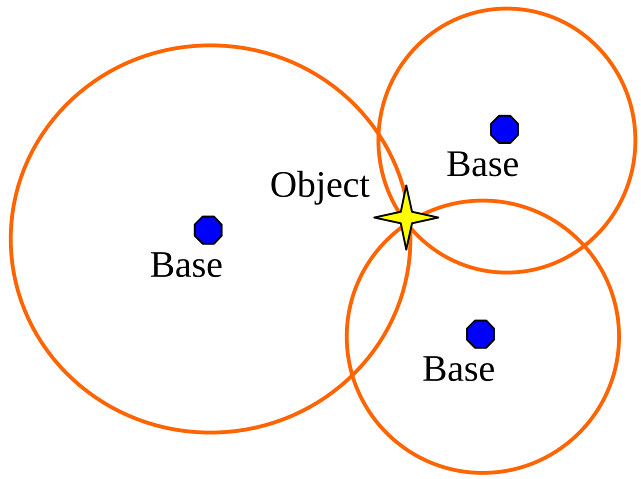

Figure 2: Location sensing concept

3.3 Implementation

Implementation of the prototype system required development of an architecture supporting the AIR ID components in a flexible manner. The following sections summarize our experience in several areas.

3.3.1 Internetworking

AIR ID basestations support only RS-232 serial connections. This is clearly a limitation since such cabling is not ubiquitous, has limited run length, and the number of physical serial ports on the server may be limited. In addition, given the need to precisely mount the AIR ID basestations in the environment, they should not depend on the proximity of the controlling server. It would be far more flexible to have basestation connectivity over arbitrary distances and configurations. The Hydra microwebserver[12] can solve this problem without any modification to the RFIDeas hardware.

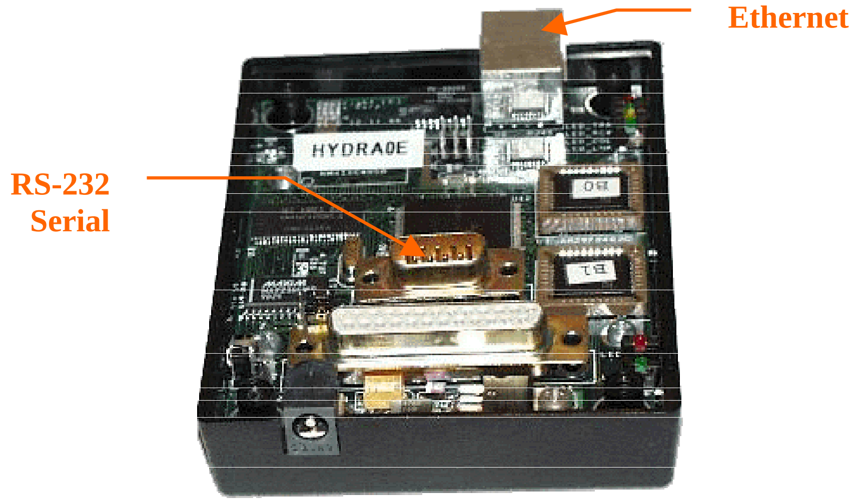

Figure 3: Hydra Microwebserver

Hydra has both an Ethernet and serial port in a small form factor and thus is ideal for the AIR ID internetworking task. A custom driver for the hydra allows pass-through operation of AIR ID commands from the Ethernet to the serial port with the addition of some caching and optimization. For example, although the AIR ID basestation protocol requires acknowledgement of each step of a command from the controlling server, the entire sequence can be sent to the hydra in one packet. The Hydra then handles the acknowledgements with AIR ID locally. The following figure depicts the overall internetworking architecture allowing a heterogeneous mix of network and serially connected basestations.

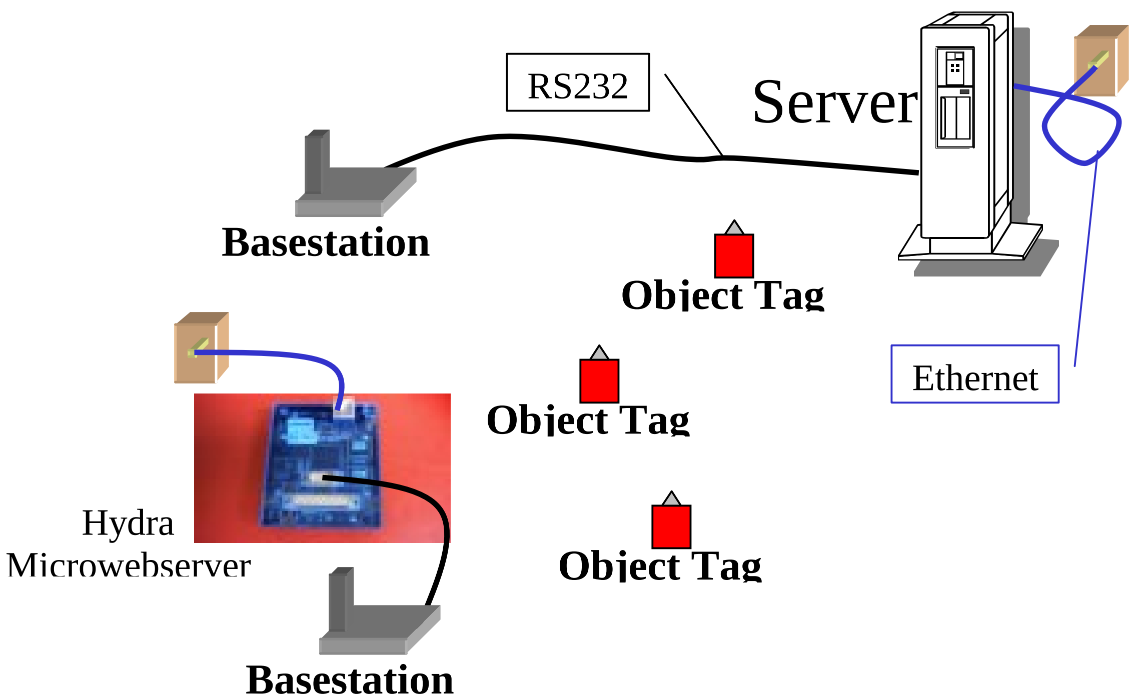

Figure 3: Architectural Illustration

3.3.2 Aggregation Algorithm

A straightforward hill-climbing algorithm attempting to minimize signal strength error relative to empirical data is used to compute the location. The algorithm works as follows:

1. Choose a random coordinate position *s* and constant distance *m*.
2. For each of the 6 possible points *p* located *m* units in the primary xyz directions, do the following:
    a. Compute a prediction vector V_p of signal strength values based on the distance *d* to each basestation from point *p* using Equation 1 derived from empirical data.
    b. Compute an error vector E_p where each element of E_p is a difference of squares of the predicted and observed signal strength values.
3. $s <= p$ where E_p is less than the current minimum error value.
4. When *s* does not change then return *s*.

Equation 1 is the function based on empirical data mapping distance to a basestation (*d*) to a signal strength estimate (*SS*). Note that SS is in abstract units.

$$SS = 0.0236 * d^2 - 0.629 * d + 4.781$$

Equation 1: Signal Strength from Distance Computation

The move distance *m* is chosen to be a constant small value mapping to approximately 6 inches of distance. An easy optimization to this algorithm is to choose *s* to start at the last computed position instead of at random with each new location computation.

3.3.3 Server

As mentioned in section 3.3.1, the AIR ID interface has been abstracted away to allow connectivity over both RS-232 and Ethernet via the hydra microwebserver. In fact, both of these protocol implementations are instances of a generic API such that it would take little effort to support other control mechanisms such as USB or a short-range radio technology. The location computation algorithm is similarly designed and it would be trivial to replace the algorithm we presented in section 3.3.2 with another method such as simulated annealing.

In this prototype system, we only publish the immediate location state of objects in the field of view for use by client applications. There is no data store or time sequence analysis. Moreover, client applications must conform to a fixed API defined by the server. While these limitations are clearly not desirable in a robust location sensing system, they are quite acceptable for this prototype in proof of concept.

3.3.4 Visualization Client

A visualization client provides display of the computed location data. This client application aided development and allows demonstration of the prototype. It is built on an OpenGL codebase developed by Zoran Popavic for a subdivision surfaces rendering system.[13] It has intuitive mouse-based navigation for pan, tilt, zoom, and 3-axis motion. Figure 4 shows a screen capture. The grid represents the physical room, the large ball is the probable position of the sensed tag, the lines are the signal strength measurements, and the small ball trail is the path taken by the convergence algorithm.

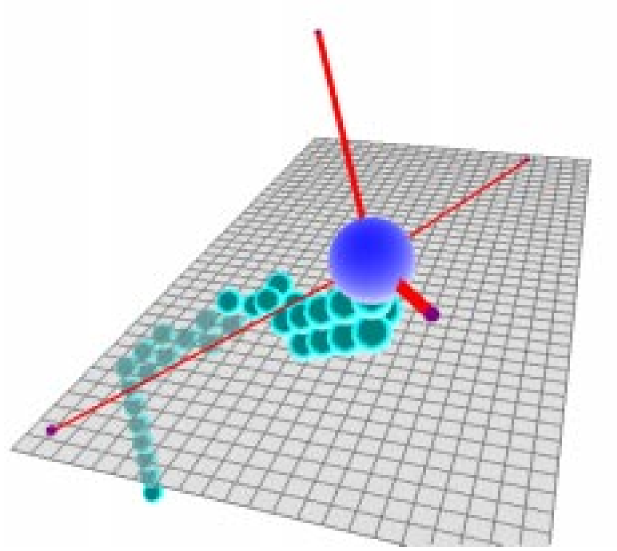

Figure 4: Visualization Environment

3.4 Results

The prototype AIR ID location system we built can determine tag location in a semi-reliable manner. However, limitations are significant. First, although the system is certainly adequate for some applications such as automatic room light control, the overall accuracy is much poorer then desired: objects can only be fixed to a position voxel about 3 meters on a side. Therefore, for small rooms, this prototype system is not much better than a rudimentary motion sensor (albeit one that could recognize the moving object by name). The coarseness is certainly a result of the mere two bits of signal strength accuracy provided by AIR ID. RFIDeas promised next generation AIR ID hardware with 8-bit accuracy but it was not available during the time of this work.

The second problem is measurement frequency. Since the AIR ID protocol is not designed for the location-sensing task, even using the API efficiently on a multithreaded server still requires 10 to 20 seconds to take one location measurement from all relevant basestations. It is very easy to miss significant changes in position of the tags as they come and go with such a slow sample rate.

Although the signal strength measurements presented in Section 3.1 were promising, one might ask why the attempt was made to build a full system using AIR ID when it may have been reasonable from the outset to conclude that two bits of accuracy (although we were working with RFIDeas for an 8 bit prototype) and a wrapped API were too limiting. The answer comes in two parts. First, other then the development labor, using the RFIDeas product is a relatively low cost solution – they are inexpensive off-the-shelf components. Second and most important, building the AIR ID based prototype gave first hand experience in the location sensing problem space and highlighted the critical design parameters. Then, using this practical experience we could construct new custom hardware much more optimized to the task.

4 New Custom Hardware

The RFIDeas prototype captured key parameters required in a well-designed location sensing system based on RF signal strength:

- It should provide a precise signal strength measurement – probably at least 8 bits of resolution.
- It should be able to attenuate transmit power in a predictable manner.
- It should allow location sampling at a rate of 1-2Hz or better.
- It should consume very little power.
- It should have sufficient memory and cycles to permit basic caching, authentication, or other such tasks.

This figure presents the new SpotON tag radio architecture we developed to address these issues.

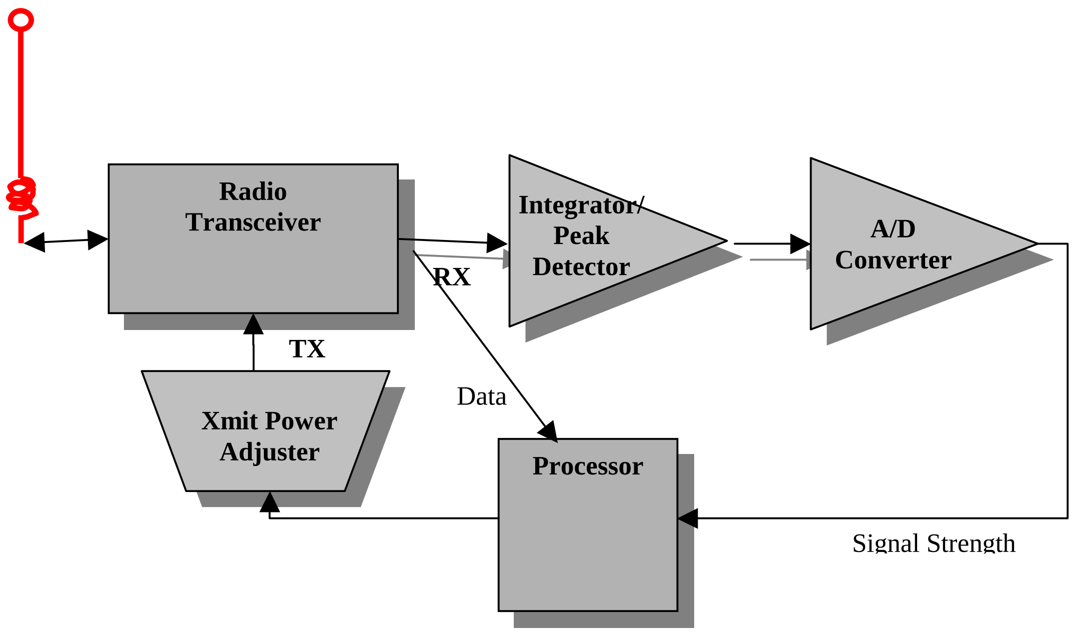

Figure 4: Hardware RF Architecture

4.1 Lab Prototype

The first step was to choose a radio transceiver and attempt to develop a method of accurately and reliably capturing the RF signal strength. Using the RFMonolithics TR1000 transceiver development boards, we prototyped several signal strength analysis circuits before settling upon one that is robust, efficient, and accurate to the degree desired. Lab tests also suggested it was reasonable to believe that a single hardware architecture could act as either a basestation or an object tag with only a difference in peripherals and the form of the packaging.

4.2 SpotON Tag Hardware Specification

Once the lab prototyping was complete, we built a custom development board. The following laundry list describes the hardware and a photo of the SpotON board is shown in Figure 5.

1. RFMonolithics 916.5MHz TR1000 Radio Transceiver
2. Custom signal strength analysis and RF attenuation circuitry
3. 10-bit A/D converter for signal strength capture
4. Dragonball EZ (68k) microprocessor
5. 2 axis MEMS accelerometer
6. 2MB DRAM, 2MB Flash
7. RS-232 port (for development only)
8. Very low power (operates on 2 Li coin cells)
9. Power Boost Regulator
10. Piezo audio speaker and 4 general purpose buttons

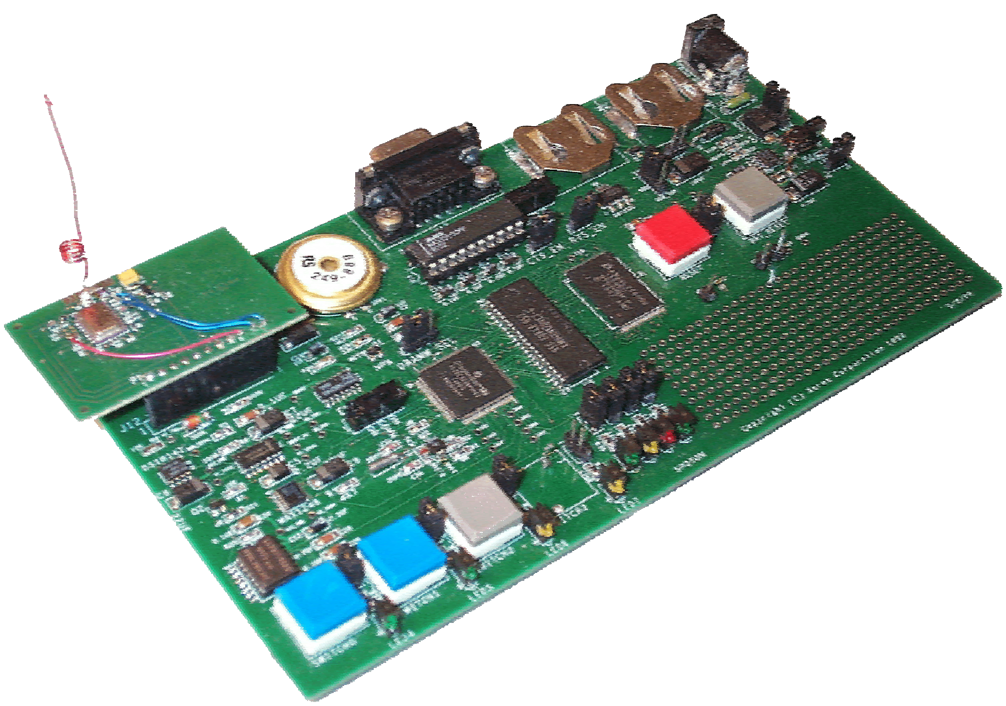

Figure 5: Custom Development Board

4.3 Cost Analysis

Figure 6 presents a breakdown of the component cost for the SpotON board. Note that the printed circuit board and its associated manufacturing costs are omitted from this chart. The total hardware component cost is around $120 but is estimate to be closer to $30-40 after revision to omit features only needed for development and assuming manufacturing in quantity.

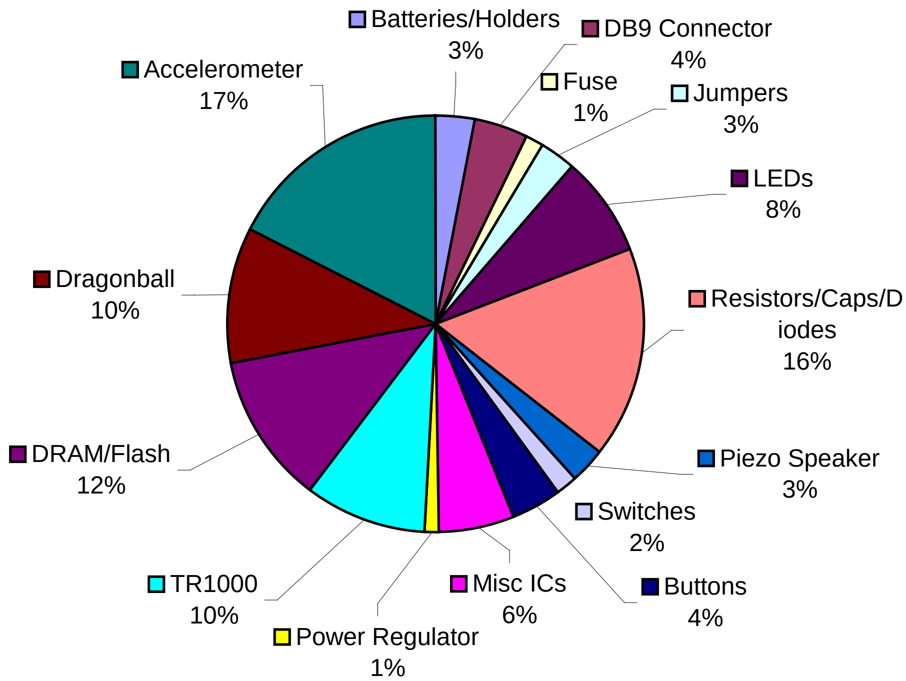

Figure 6: Cost Breakdown of the SpotON Board

4.4 Power Consumption

A basic power consumption analysis on the SpotON board has been completed to quantify the power requirements of various components and operations. The results are presented in Figure 7. Consider these numbers as upper bounds because the development board is optimized for debugging and not power efficiency. In the present design, two lithium coin cell batteries provide approximately 180mAh each implying about 10 hours of normal operation. Even with careful power saving techniques, this battery life is less than desirable. Therefore, the next hardware revision may increase the number of coin cells or move to larger cells such as AAAA.

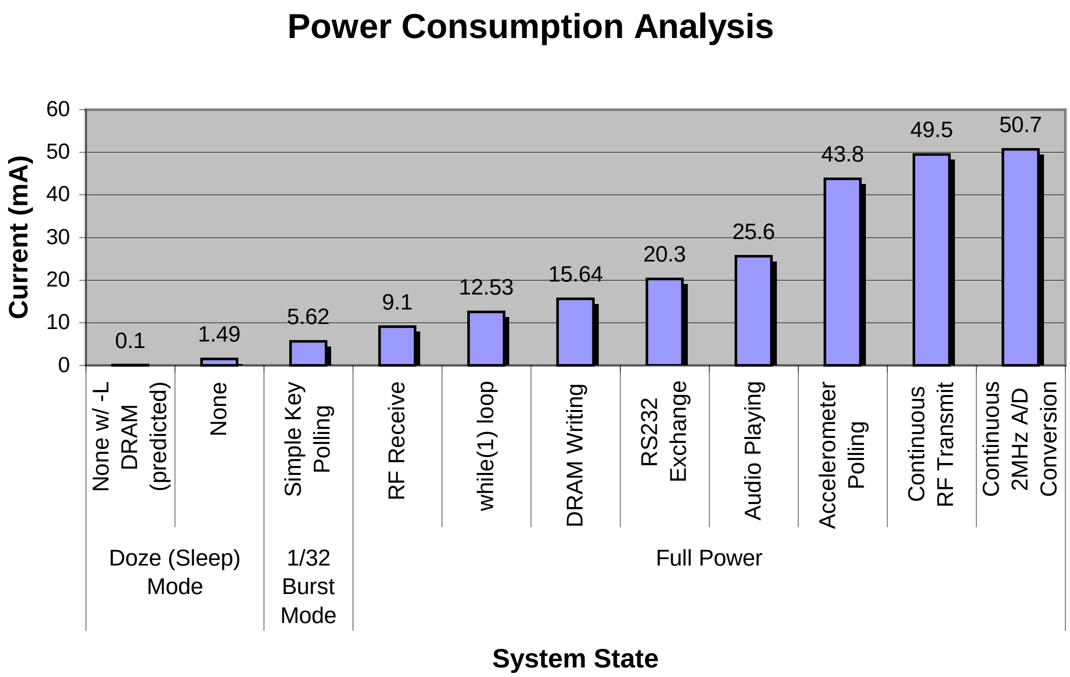

Figure 7: Power Consumption of the SpotON Board

4.5 Status and Goals

The firmware is nearly complete. A hardware revision to remove features only used for development, make minor design adjustments, and shrink the board to tag size will occur shortly. One possible idea is to implement an interface and form factor so that SpotON tags can function in a Compact Flash, PCMCIA, or a similar generic card slot allowing many devices to plug in SpotON and easily integrate location-sensing ability. In any case, SpotON tags would certainly also work standalone. Lab experiments with the new hardware seem to indicate that sub 1^3 meter voxel location sensing accuracy should be possible but since the actual development is still in progress, there are no hard numbers to back up this claim at present.

5 Applications

Fine-grained location information would enable applications not readily possible by other means. Although this is not an applications paper per se, we can nonetheless briefly speculate about a few possibilities. Precise positioning info could be used to dynamically track people and mobile objects in real time within a space such as by a camera, lighting instrument, or laser pointer. This could be valuable for surveying security sensitive articles or perhaps by a multi-camera cinematographic video conferencing system. SpotON tags may also be useful in much more esoteric applications such as presentations, film, or theater production. One could imagine a historical marker system by which interesting information or anecdotes could be "placed" around a site and retrieved dynamically as location aware objects and people travel through a space. For home automation, various consumer electronics and household appliances could present their interface or take action based on the location of people or other tagged objects. The application we are currently developing at the University of Washington is one for migrating mediated audio. Location-sensing technology such as SpotON tags allow multimedia streams to follow roaming users through different media-cells. The content is mediated when multiple users come together in the same space.

6 Future Work

Although developing the SpotON tagging technology for fine-grained location sensing is an important contribution, location sensing as a field is ultimately about data integration and sensor fusion. Although a system such as SpotON may provide very accurate location information of a particular type, it alone is almost certainly not the ultimate solution in the problem space. Any robust and scalable location sensing architecture needs to develop a data model to characterize the notion of location as described by a heterogeneous mix of sensors. For example, a complete system may use SpotON tagging technology supplemented with the onboard accelerometers, wide area GPS receivers, and perhaps a simple software agent to determine when a person is typing on a keyboard. Here is a brief taxonomy of several location data types that should be successfully representable in a general location model:

1. Absolute – Systems that provide conclusive position data in the form of coordinates, latitude and longitude, etc. SpotON (properly configured) and the GPS satellites are two examples.
2. Relative – Location data generated from a relative sensor is comparative to another sensor but it alone does not yield an absolute location measurement. Many systems including SpotON can be used to provide relative location data instead of or in addition to absolute positions.
3. Anonymous – Traditional motions sensors produce anonymous data, as does the keyboard agent previously mentioned (at least in the basic case). The sensor is aware that someone or something is present but their ID is unknown. Anonymous data is also generated if the user of a positioning system chooses not to reveal his or her identity to the infrastructure in the interest of privacy or security.
4. Symbolic – This type of data does not have a geographic form, although it can often be mapped to such a form with another step. The canonical example is the event indicating someone is typing at a particular keyboard and the machine name or IP address is known but its precise geographic position is not.
5. Motion – The key ingredient in motion data is the addition of the fourth dimension: time. Accelerometers and radar speed measurement systems both provide direct measurement of motion data such as speed and acceleration. Motion data can also often be derived by aggregating other data such as relative position information over time. Motion is arguably an orthogonal issue and perhaps not a distinct data type, but we include it in this list anyway because it is certainly an essential part of a robust location-sensing system.

Assuming such a general location representation model can be developed, several important issues then need to be addressed. For example: Which sensor fusion techniques are best suited to populate the model in different circumstances? How do concerns such as grain-size variability or inconsistency from different location sensing sources influence the model and data-store design? How do privacy and security concerns associated with exploiting location information in applications influence the model? How can location-sensing data be best managed and integrated into a larger context-aware architecture or is this not necessary? Finally, what applications do real users want from a location sensing system?

7 Conclusions

This paper has presented the background and development of SpotON – a fine-grained indoor location sensing system based on RF signal strength. SpotON tags are custom devices that operate standalone or potentially as a plug in card enabling larger devices to take advantage of location-sensing technology. They are low power, small form factor, and can be located with sub $1m^3$ voxel accuracy, yet still have enough processing power for caching, authentication, and other such tasks. We believe that SpotON tagging technology offers several advantages over existing systems and, when combined with a general location data model, will be a key component in the next generation of calm and invisible computing.

8 Acknowledgements

The work presented in this paper was completed over the course of summer 1999 working with Roy Want, Marshall Bern, and Frank Vest at Xerox PARC and during autumn 1999 and winter 2000 at the University of Washington working with Gaetano Borriello and the Portolano research group. Thanks go to all of these folks for their valuable insight and assistance.

---

[1] Starlink Incorporated. “DGPS Explained,” http://www.starlinkdgps.com/dgpsexp.htm. 1999
[2] R. Want, A. Hopper, V. Falco, J. Gibbons. “The Active Badge Location System,” *ACM Transactions on Information Systems*, Vol. 40, No. 1, pp. 91-102, January 1992
[3] R.Want, B.Schilit, N. Adams, R. Gold, D. Goldberg, K. Petersen, J.Ellis, M. Weiser. “The Parctab Ubiquitous Computing Experiment”, Book Chapter: “Mobile Computing”, Kluwer Publishing, Edited by Tomasz Imielinski, Chapter 2, pp 45-101, ISBN 0-7923-9697-9, February 1997. Also available as Xerox B&W, March 1995 P95-00050.
[4] Microsoft Research. “Easy Living”, http://www.research.microsoft.com/easyliving/, 1999
[5] M. Turk, “Moving from GUIs to PUIs,” *Proc. Fourth Symposium on Intelligent Information Media*, Tokyo, Japan, December 1998. (also Microsoft Research Technical Report #MSR-TR-98-69)
[6] P. Bahl, V. N. Padmanabhan. “User Location and Tracking in an In-Building Radio Network”, *Microsoft Research Technical Report*: MSR-TR-99-12 (February 1999)
[7] PinpointCO Website, http://www.pinpointco.com/, 1999
[8] A. Harter and A. Hopper, “A Distributed Location System for the Active Office.” *IEEE Network*, January 1994
[9] Federal Communications Commission, “Enhanced 911”, http://www.fcc.gov/e911/
[10] Bible, S.R., Zyda, M. and Brutzman, D. (1995). “Using Spread-Spectrum Ranging Techniques for Position Tracking in a Virtual Environment” in the Proceedings of the 1995 Workshop on Networked Realities, Boston, MA, October 26-28, 1995.
[11] RFIDeas Website, http://www.rfideas.com/, 1999

[12] J.H. Kang, K. Partridge, "Hydra Microwebserver", http://www.cs.washington.edu/research/portolano/hydra/, 1999

[13] Zoran Popovic, "Motion Transformation by Physically Based Spacetime Optimization." *Ph.D. Thesis*, Carnegie Mellon University 1999.