PIP Tags: Hardware Design and Power Optimization

Bernhard Firner, Shweta Medhekar, Yanyong Zhang, Richard Howard, Wade Trappe
WINLAB, Rutgers University
North Brunswick, NJ 08902, USA

Peter Wolniansky
Radio Technology Systems
82 Mt Tabor Way
Ocean Grove, NJ 07756
wolniansky@optonline.net

Eitan Fenson
Inpoint
770 Lexingtion Ave, Sixth Floor
New York, NY 10021
e-fenson@inpointsys.com

{bfirner,shwetam,yyzhang,reh,trappe}@winlab.rutgers.edu

Abstract

Currently, it is difficult to find radio devices that meet the size and lifetime requirements of many continuous tracking applications. This paper will discuss the requirements of these applications, and then describe our experiences in designing a hardware platform that satisfies these requirements. We focus on the steps needed to reduce power consumption in radio devices. This paper can also serve as a step by step tutorial for future attempts to minimize power consumption in similar platforms.

Categories and Subject Descriptors

B.8.2 [**Hardware**]: Performance and Reliability—*Performance Analysis and Design Aids*

General Terms

Design, Measurement

Keywords

Active RFID, Asset Tracking, Power Efficiency

1 Introduction

RFID technology sits at the forefront of the future of pervasive computing in the business world. The potential of RFID embodies the ideas of pervasive computing - a constant flow of information for every item from the moment it is manufactured until the time it is finally recycled. Integration of RFID technology into supply chain management would help to reduce theft, would decrease labor costs for inventory tracking, and would increase product management efficiency by reducing loss of sales due to out of stock issues. The potential exists for applications beyond that - integration of RFID technologies into hospitals could allow staff to track equipment, patient charts, and medicine. Quickly locating such items can be vital if the time spent locating all the necessary items to start an operation takes as long as the operation itself.

However, there are many barriers to the adoption of RFID as a complete replacement for a mature tracking technology, such as bar codes, from both business and technology standpoints [15]. The slow and difficult deployment that Wal-Mart and many of its suppliers have seen after Wal-Mart's RFID mandate suggest that RFID technology requires further development to be useful in the business supply chain. The problems with passive RFID tags - poor range, difficulty avoiding transmission collisions in dense environments, and poor performance when attached to certain items - can be overcome with semi-passive or active tags. Active tags have better performance than passive tags because they have more energy available for the radio link, but they have limited lifetimes associated with the battery, are much more costly, and are much more bulky than passive tags.

We have developed a new active RFID system named Roll-Call™ that will solve the cost, lifetime, and size issues of currently available active RFID systems. This paper will focus upon the hardware development of the RFID tags for our system, with a special emphasis on the steps taken to minimize the power consumption of the tags. The lessons we learned in this exercise can be applied to any sensor network application that requires energy reduction at the node level.

This paper will first analyze the related devices in this area and explain the need for a new active RFID design in section Section 2. Our choice of hardware architecture and design is presented in Section 3 and the steps we took to minimize the power consumption of our tags is described in Section 4.

2 Roll-Call™ Architecture

The Roll-Call™ system targets a family of continuous, real-time asset tracking applications with many challenging requirements. The requirements that specifically affect the hardware design of the tag are:

- *Low Cost:* The total cost of an RFID system can be calculated as
$$C = N_{tag}C_{tag} + N_{reader}C_{reader},$$
where the number of tags, N_{tag}, is determined by the number of asset items in the inventory, and the number of readers, N_{reader}, is governed by the number of tags that can be handled by a reader and the desired overlap degree between readers. The envisioned applications

HotEmNets'08, June 2–3, 2008, Charlottesville, Virginia, USA

RFID tags	Dimensions (xcm^3)	Density (per reader)	Tag Cost ($)	Reader Cost ($)	Lifetime (years)
Smartcode[9]	diameter of 4cm	hundreds	> 50	> 50	N/A
AeroScout[1]	$6.2 \times 4.0 \times 1.7$	hundreds	> 50	> 50	2.33 at a 30-sec beacon interval
PanGo[5]	$6.4 \times 4.3 \times 1.8$	hundreds	> 50	> 50	5+
Axcess[2]	$7.1 \times 3.1 \times 0.7$	hundreds	34	545 - 1099	2-3
RF Code[8]	$4.7 \times 3.4 \times 1.2$	hundreds	> 50	> 50	7 at a 10-sec beacon interval

Table 1. **A comparison of the state-of-the-art active RFID systems.**

will have $N_{tag} > 1000$. As a result, the tags and readers must be made low-cost, and in particular, each tag should ideally cost less than a dollar in high enough volume.

- *Short Beacon Interval:* Continuous asset tracking requires each RFID tag frequently report to the reader, in order to minimize report latency during theft or destruction of an item. A beacon interval of one second or less is preferable for real-time tracking.
- *Small Size:* The tags should be small enough that they do not obstruct observation, handling, and storage of the item they are attached to.
- *Long Lifetime:* Even if the tags transmit a beacon every second the system should still operate over a meaningful period of time, on the order of a year or more. The previous requirement, small tag size, makes this more difficult since a large battery cannot be used.

Keeping the above requirements in mind, we try to find suitable RFID systems in today's market. Today's RFID tags can be broadly categorized into two classes: passive RFIDs and active RFIDs. Passive tag systems, such as those from Tagsys [10] and Paxar [6], cannot provide continuous tracking due to the range limitations (about a meter) and small densities supported by readers in passive tag systems (a reader usually supports a few tens of passive tags per read).

A variety of active tags exist on the market. Table 1 summarizes several representative active RFID tags. Though many of these tags are designed for asset tracking, they are unsuitable for *continuous, item-level* asset tracking. First of all, the tags are costly (more than $50). Second of all, many tags are relatively large in size, inappropriate for small items such as jewelry. Third of all, the lifetime of many tags is not long enough if we scale them by the desirable beacon interval. For example, a 2.33-year lifetime at a 30-second beacon interval will have a much shorter lifetime if the beacon interval was instead one second. The exact change in lifetime is dependent on what portion of the energy is used during transmission but, since transmitting consumes much more energy than the idle state, increasing the number of transmissions will have a dramatic impact on the lifetime of a tag. For these reasons it is imperative to design a new system to satisfy these requirements.

Our system, Roll-Call^{TM}, is designed to overcome the limitations of these other systems. Since we have kept the cost and effectiveness of the tag and basestation hardware in mind from the beginning, an emphasis has been placed on having the software backend compensate for low complexity hardware design. The data flow in the system is all unidirectional as shown in Figure 1 - the tags are never required to respond to any signals. Packets are received by the basestations which then forward the data to a processing station running on a PC. The processing station processes the data and stores it in a database where it is accessed by a web-based front end. This saves energy in the tags since they never need to receive. The "dumb" tags are compensated for with backend processing. Since the exact content of each packet is known the packets involved in a collision can be repaired with signal processing of the received signal. A very small transmitted signal also decreases the likelihood of collision. As a result, the adverse impact of collisions is less severe than it it may at first appear.

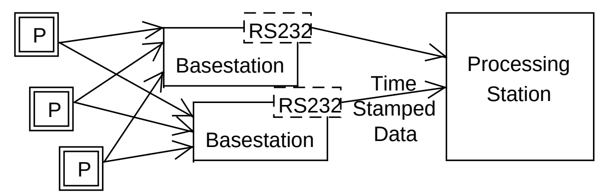

Figure 1. **Dataflow Diagram.**

The higher complexity of the back end allows the hardware design to be simplified and for costs to be held to a minimum. The only function of the tags will be to periodically announce their presence by transmitting a short ID message. All other information about the item will be stored in a network database. Physically, tags must be kept as small as possible while having a lifetime on the order of a year or more. In reference to their small size and the small amount of radio traffic they generate we have named the tags "pipsqueaks", or "PIPs" for short. Of course additional effort must be spent making the inexpensive radio tag capable of supporting the target application, especially for the lifetime of the tag. That effort is the focus of this paper and will be discussed in detail in the rest of this paper.

3 Tag Hardware Design

During every step of the design process we focused upon the end goal of an affordable final design when executed in the lowest cost custom hardware. To allow for flexible prototyping, we chose to adopt off-the-shelf commercial embedded chipsets that are low-cost, low-power, easy to program, and with a small form factor. Ultimately we will need only a single chip with a simple radio and some logic for loading the ID into the radio. With these objectives in mind, we chose the TI chip that includes a 16-bit Silicon Laboratories C8051F321 microprocessor and a Chipcon CC1100 radio transceiver. The tags and basestations share the same

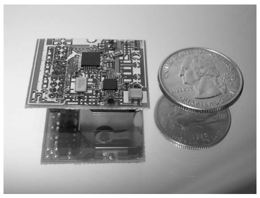

(a)

(b)

Figure 2. **(a) PIP. (b) Base-station: Texas Instruments CC1100/CC1150-868/915 MHz Development Kit**

chipset and radio and are both programmable in C. The operating voltage of the chip is listed as 2.7 to 3.6 volts and the temperature range is -40 to 85°[14]. The tag runs on a low cost, high energy density 20 mm diameter lithium coin cell, the CR2032 battery[4].

Although they are not used in this project, the tag has an onboard temperature sensor and has pins available for attaching additional sensors. The board also has a USB connection available and a small assortment of LEDs.

The microcontroller supports a proprietary instruction set from Silicon Labs called MCS-51^{TM}. It can be programmed with standard 803x and 805x assembly or with a provided C compiler. The microcontroller has access to 1280 bytes of data RAM and 16kB of reprogrammable FLASH program memory. The maximum clock rate of the microcontroller is 25MHz, but this is adjustable.

The Chipcon CC1100 radio chip can be programmed to operate on bands starting at 315, 433, 868, and 915 MHz and claims a sensitivity down to -110 dBm at a data rate of 1.2 kbps[14] with a 1% error rate. Supported modulations are OOK, ASK, FSK, GFSK, and MSK. Transmit power can be adjusted from -30 dBm to +10 dBm. The entire radio fits on a 4x4mm package. The device is a transceiver, but *we disabled the receiving function on the tags to conserve energy*. We do use the receive function in the basestation, which uses the same tag for both RF and digital processing.

The PIP has an antenna built into the chip. Since the chip is not large enough to house a 900MHz dipole antenna, an elecrically small antenna follows part of the circumference of the chip. Though the performance of this antenna is inferior to a full-size antenna, the compromise was necessary to save space and slim the tag down to its current size of 3cm x 2.5cm x 0.5cm, with the height mostly due to the size of the battery. The antenna inefficiency is compensated for by the extra power made available for transmission, and its low Q makes the tuning insensitive to the proximity of metal. This may seem large, but part of the size is due to the programming pins and circuit layout. The actual size of a manufactured PIP would be constrained by the size of the battery. At 20mm it is comparable in size to the bar code tag currently required on jewelry (although heavier).

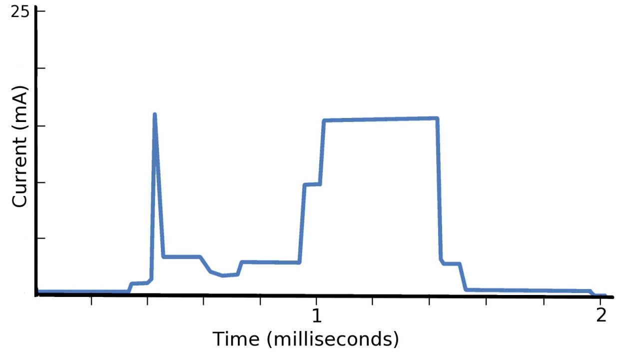

(a)

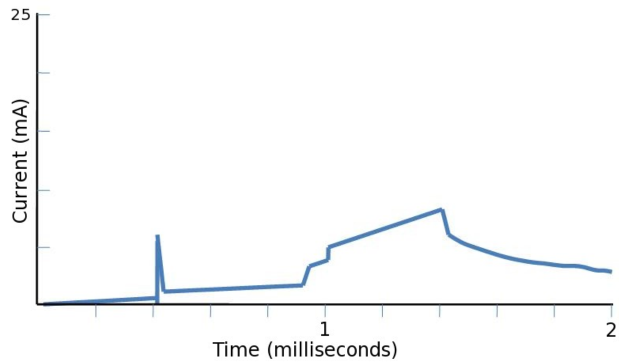

(b)

Figure 3. **The current drawn from the battery during the wake up and transmission part of the duty cycle, before and after the addition of a 100 microfarad chip capacitor.**

The CR2032 costs less than $0.14 even in relatively low volume (4,000)[3]. The nominal energy density of the cell is about 190 mA-hours down to a voltage of 2.7 volts, the lower operating voltage of the microcontroller that we are using. These specifications are provided at a standard load current of only 0.2 mA, though, and during the transmit phase our tags draw a maximum of about 16 mA for about 400 microseconds (0 dbm and 168 bit total packet length). We manage this issue by including a 100 microfarad chip capacitor on the tag to provide the pulse current and level out the battery load. As can be seen in Figure 3, the addition of a capacitor smoothes the large energy spikes over a longer period of time. Three months of continuous monitoring show that the altered power consumption during the transmission phase results in battery aging consistently with the recommended current load (0.2 mA) and we have not observed a significant reduction of expected energy density. It is worthwhile to note that there is active work in the battery industry to provide cells with higher pulse current density for sensor applications like ours and we expect performance to be better characterized and improved in the near future[7].

The basestations use the same hardware as the PIPs. Since they are not under the same space constraints, they have a tuned 900MHz dipole antenna attached. The current basestations also include an RS232 serial port running at 115200 bps, or an Ethernet or WiFi connection constrained to the same speed, for data transfer. By choosing this hardware, costs were held to a reasonable level for the prototyping phase of our project. Even after settling upon a particular system we still needed to make many modifications before it was capable of supporting our requirements, mainly due to the long lifetime and low power requirements, similar to what has been seen in systems such as [13].

4 Energy Minimization - Mining the Joules and Bits

The first step in minimizing energy draw is obvious - all unused devices and sensors should be turned off. The PIPs use a crossbar network to connect the output and input pins of sensors, LEDs, and other peripherals (including the radio). All unused devices were disabled in the crossbar.

Initial tag testing began by setting the beacon transmit interval to one second with a 208-bit payload using MSK modulation and a transmit power of 10 *dBm*. Using one CR2032 battery with a capacity of 190 *mA-hr*, the initial system only ran for 12 hours before draining the battery. After conducting an in-depth investigation on the tag's operation, we classified the tag's energy usage into the the following five regions: (A) the waiting period between two beacons, (B) the microcontroller transfer phase (to radio), (C) the radio initialization phase, (D) the radio transmission phase, and (E) the radio calibration phase. Among them, region *A* has the longest duration (just less than a second) and, in the initial setting, region *A* consumes almost all of the energy. Thus region *A* will be our first target for energy minimization.

4.1 Optimizing Region A

The default hardware configuration leaves both the microcontroller and radio active during region *A*, resulting in unnecessary power consumption. As was discussed previously unused peripherals were turned off in the PIP. However, during the idle state *all* of the peripherals, including the radio, are not in use and can be turned off. In platforms where the radio is used for receiving this is a step than cannot be performed, although wake on radio features may be used. A clock is needed bu the 25MHz microcontroller is unnecessary - a much slower clock rate would be adequate.

Since the duration of region *A* is independent of the clock speed it makes sense to use a slow clock during this idle period to conserve energy. Since no computations occur in region *A* the use of a slower clock has no disadvantages. For an 8051 microcontroller system, however, there is a limit on the maximum value of capacitance (*C*) and resistance (*R*) to be connected in an *RC* circuit to have a stable, temperature independent oscillator. The minimum recommended value is 8 *kHz*. Reducing the clock speed to 8 *kHz*, compared to the original 6 *MHz* system clock, during sleep time significantly reduces the energy consumption in region *A*.

Previously, the radio stays on in region A, and at the end of the 1 second period, the microcontroller must transfer data to the radio. Changing the clock rate and shutting down the radio will change the tag's behavior slightly. To take advantage of the slower clock, the tag switches to the 8 *kHz* clock after entering region *A* and registers an interrupt that will trigger after one second and force the tag into region *B*. When the interrupt triggers, the tag switches to the faster clock (to ensure fast data transmission), turns on the radio, and transmits a packet. These optimizations reduced the energy in region *A* from 49800 $\mu J/pulse$ with the radio on and microcontroller running at full speed to 72 $\mu J/pulse$ with the radio off and a slow microcontroller clock.

SYSCLK (MHz)	SCLK (MHz)	TIME (ms)	CURRENT (ms)	ENERGY ($\mu J/pulse$)
1.5	0.75	1.43	3	12.87
3	1.5	0.865	3.5	9.083
6	**3**	**0.575**	**4.25**	**7.32**
12	6	0.44	7	9.24
24	6	0.385	16	18.48

Table 2. Energy in Region B with different SYSCLK's and the best SCLK.

4.2 Optimizing Region B

In region *B*, the microcontroller transfers the data to the radio using the on-board 4-wire serial interface. The energy consumed by this region depends upon the speed of the microcontroller clock (*SYSCLK*), the serial clock (*SCLK*) (which is internally generated by passing the *SYSCLK* through dividers to pace the data transfer) and the amount of data to be transmitted. Fixing the amount of data to be transmitted, we can tune the frequencies of *SYSCLK* and *SCLK* to achieve power efficiency. Increasing the clock speeds reduces the time of internal data transfer and initialization at the expense of higher current draw during the period. Hence, clock speed and current have to be carefully balanced to minimize the overall energy consumption.

We empirically tuned the values of *SYSCLK* and *SCLK* to achieve the minimum energy consumption. In this process, we first varied the value of *SYSCLK*, and for each *SYSCLK* frequency, we next varied the *SCLK* value to search for the best combination that results in the lowest energy consumption. We note that the maximum frequency allowed for *SCLK* is half of the *SYSCLK* frequency up to a maximum of 6 *MHz* [14][12]. We report several typical *SYSCLK* values, as well as the correspondingly best *SCLK* values in Table 2. From the table, we observe that for a given *SYSCLK*, the maximum allowable *SCLK* always gives the least energy consumption. The best configuration across the entire spectrum occurs by setting *SYSCLK* to 6 *MHz* and *SCLK* to 3 *MHz*. In these experiments, the duration of region B (presented by the TIME column in Table 2) and the current drawn in the region (presented by the CURRENT column) were measured by observing the time-dependent current drain of the PIP on an oscilloscope.

4.3 Optimizing Region C

Region C is the time taken to initialize the radio. This is a constant 80 μs, irrespective of system clocks and the size of data. This region is not user controllable.

4.4 Optimizing Region D

In region D, the actual data transmission takes place. The energy spent in region D is determined by the three factors: (1) the amount of data to be transmitted, (2) the data rate, and (3) the *RF* power of transmission. Among these three factors, the first two directly impact the duration of region D while the third factor directly impacts the power level. As a result, in order to reduce the total energy consumption in region D, we should reduce the data to be transmitted, increase the data rate, and decrease the transmission power level.

The amount of data to be transmitted is usually controlled by either the application or the RFID specifications, and might be unchangeable. However, since we are designing our system with the hardware in mind we can choose to minimize the packet size by only sending a portion of the entire tag ID with each packet. The portion of the ID being transmitted is cycled, thus reducing the packet size.

In order to maximize the PIP's effectiveness, we choose 0 dBm as the RF transmission power and 250 kbps as the radio data rate because a higher data rate may be unreliable. A transmit power of 10 dBm is actually the most power effective setting - 0 dBm consumes half the electrical power but has less than half of the over the air power as 10 dBm. This is because the radio performs optimaly at its highest settings. However, 0 dBm is sufficient for most deployment needs although a transmit power of 10 dBm might be more suitable in some environments.

4.5 Optimizing Region E

In region E, the radio calibration takes place. The time taken to calibrate the radio is constant and cannot be controlled. However, radio recalibration is only necessary to prevent radio drift as the operating characteristics change due to temperature or operating voltage. The target environments for this system will not undergo sudden temperature shifts though, so doing this every duty cycle is unnecessary. Radio calibration can be done infrequently and the energy cost of this region will be amortized over many duty cycles. With such a configuration, the energy of region E can be ignored for the purpose of estimating battery lifetime.

4.6 Further Optimizations

Even with all devices and inputs disabled in the crossbar network high power consumption was still observed with changes in humidity and board cleanliness. Care must be taken with the circuit board layout and during packaging to avoid sources of potential surface leakage. Such flaws could draw tens of μamps of power, significantly reducing the lifetime of a tag with such a limited power source Cleaning the well-handled test devices reduced the observed current in some areas and removed artifacts due to surface leakage.

One last energy-saving optimization was done by minimizing the radio link error. We observe that merely matching the carrier frequencies of the *PIP* and the basestation does not guarantee the most reliable radio operations due to discrete quantizations of the radio frequency. To correct for this, we used the transmission success as feedback to tune the basestation. The basestation frequency was initially set to the same frequency as the PIP and was then increased in steps of 50 kHz until no more messages could be received. The process was then repeated, this time lowering the frequency by 50 kHz at each step. The center of the two extremes is then the center of the minimum-error band where the basestation should listen.

4.7 Final Configuration

The power draw in each region is shown in Table 3. With one second transmit intervals each duty cycle consumes approximately $0.031mA - seconds$ if calibration is done and $0.025mA - seconds$ without calibration. Since calibration will be done infrequently the second value will be used to estimate the total lifetime. Since the CR2032 coin cell battery has $190mA - hours$ it will support approximately 318 days of operation, a tremendous improvement over the initial 12 hour lifetime. Increasing the beacon interval could extend the lifetime further - with 10 second intervals it becomes 469 days. Amdahl's law tells us that as the power consumed during transmission becomes a smaller percentage of the total power draw per duty cycle the extra lifetime gained by increasing the beacon interval diminishes, but the lifetime can be extended in applications where a one second beacon interval is more than adequate. Future hardware changes will yield additional improvements. For instance, replacing the microcontroller with simpler circuitry and using a custom radio will further reduce the size and cost of the components. With the current hardware design the cost of manufacturing a tag is in the $5 range.

REGION	TIME	CURRENT	ENERGY	%TOTAL
A	1 s	16 μA	48μJ	64%
B	600 μs	3 mA	5.4μJ	7%
C	70 μs	10 mA	2.1μJ	2%
D	400 μs	16 mA	19.2μJ	26%
E	720 μs	8.5 mA	18.4μJ	0%

Table 3. Approximate energy consumption for each region after all optimizations using 1 second beacon intervals.

5 Concluding Remarks

In this paper we described the need for a new hardware design to support asset tracking applications. The hardware design of the Roll-Call^{TM} was described and the steps taken to create a cost-feasible solution were shown. General energy minimization techniques that can be applied to the development of any small sensor system were discussed in detail in the hope that they will provide a framework for future development of energy conscious systems.

6 References

[1] Aeroscout. http://www.aeroscout.com/.
[2] Axcess inc. http://www.axcessinc.com/.
[3] Digi-key. http://www.digikey.com/.
[4] Panasonic. http://www.panasonic.com/.
[5] Pango. http://www.pango.com/.
[6] Paxar. http://www.paxar.com/.
[7] Powermems. http://www.powermems.com/.
[8] Rfcode. http://www.rfcode.com/.
[9] Smartcode. http://www.smartcodecorp.com/.
[10] Tagsys. http://www.tagsysrfid.com/.
[11] M. Beigl, A. Krohn, T. Riedel, T. Zimmer, C. Decker, and M. Isomura. The upart experience: building a wireless sensor network. *Information Processing in Sensor Networks, 2006. IPSN 2006. The Fifth International Conference on*, pages 366–373, 19-21 April 2006.
[12] Silicon Laboratories. *C8051F320/1*. http://www.silabs.com/ mcu/.
[13] I. Talzi, A. Hasler, S. Gruber, and C. Tschudin. Permasense: investigating permafrost with a wsn in the swiss alps. In *EmNets '07: Proceedings of the 4th workshop on Embedded networked sensors.*, pages 8–12, New York, NY, USA, 2007. ACM.
[14] Texas Instruments Norway AS. *CC1100 Single Chip Low Cost Low Power RF Transceiver*, June 2006. Data Sheet (Rev.1.1).
[15] A. W. L. Wong and D. Viehand. Critical Issues that Will Determine the Future of RFID. *Management of Mobile Business, 2007. ICMB 2007. International Conference on the*, page 48, July 2007.