What Am I Looking At? Low-Power Radio-Optical Beacons for In-View Recognition on Smart-Glass

Ashwin Ashok, Chenren Xu, Tam Vu, Marco Gruteser, Rich Howard, Yanyong Zhang, Narayan Mandayam, Wenjia Yuan, and Kristin Dana

Abstract—Applications on wearable personal imaging devices, or *Smart-glasses* as they are called, can largely benefit from accurate and energy-efficient recognition of objects that are within the user's view. Existing solutions such as optical or computer vision approaches are too energy intensive, while low-power active radio tags suffer from imprecise orientation estimates. To address this challenge, this paper presents the design, implementation, and evaluation of a radio-optical hybrid system where a radio-optical transmitter, or *tag*, whose radio-optical *beacons* are used for accurate relative orientation tracking of tagged objects by a wearable radio-optical receiver. A low-power radio link that conveys identity is used to reduce the battery drain by synchronizing the radio-optical transmitter and receiver so that extremely short optical (infrared) pulses are sufficient for orientation (angle and distance) estimation. Through extensive experiments with our prototype we show that our system can achieve orientation estimates with 1-to-2 degree accuracy and within 40 cm ranging error, with a maximum range of 9 m in typical indoor use cases. With a tag and receiver battery power consumption of 81 μW and 90 mW, respectively, our radio-optical tags and receiver are at least 1.5$\times$ energy efficient than prior works in this space.

Index Terms—Smart-glass, low-power, positioning, recognition, RFID, optical, infrared, wearables, tags

1 Introduction

Smart-glass or wearable personal imaging devices provide endless possibilities for applications that can interact with the physical world. Today, there are multiple smart-glass devices commercially available in the market enabling plethora of interactive applications ranging from visualizing content on a smart-glass heads-up display [1], [2] to virtually interacting with objects in physical space [3], [4].

In applications that interact with objects (or contexts) in the physical world using smart-glasses, the *in-view recognition problem* becomes relevant. As illustrated in Fig. 1 the in-view recognition problem on smart-glasses translates to determining the precise identity (*what is the object in the user's view?*) and relative position of the user to the object of interest (*what orientation is the object in the user's view and how far is it from the user?*). For example, smart-glasses that can recognize display items in day-to-day lives can provide more information as well as navigation to those artifacts in real-time; users can get more information about the party attendees and find whether they have any social connections with them [5], enhancing human-to-human interactions. In general, knowing precise orientation of a smart-glass user to a context and its identity also benefits a diverse set of applications such as smart advertising, gaming [6], and even for tracking user shopping behavior in stores [7]. However, the in-view recognition problem on smart-glass comes with its own set of challenges.

Low-power challenge for in-view recognition. In addition to the fact that a solution for the in-view recognition problem requires precise identification and positioning of objects, a fundamental challenge on smart-glasses is minimizing battery power usage. Interactive applications on smart-glass typically require continuous interaction with the physical space thus requiring to operate the recognition modules for long durations resulting in significant battery drain. To address this low-power in-view recognition challenge for smart-glass, in this paper, we aim to design a low-power solution for continuous and accurate in-view recognition.

Technical approaches for in-view recognition can be broadly categorized into two categories: (1) passive and (2) active. Techniques that are positioning/tracking based or computer-vision based come under the passive category. Designs, so-far in this space, typically trade-off between accuracy and battery-lifetime. For example, the Wikitude World Browser [8] adopts the positioning approach where by it uses the GPS position and compass together with map information to infer what landmarks a smartphone is pointed at. This approach is generally limited to outdoor use and the accuracy drops significantly when objects of interest are placed closely together. Computer vision based solutions

- *A. Ashok is with the Department of Electrical and Computer Engineering, Carnegie Mellon University, Pittsburgh, PA 15213. E-mail: ashwinashok@cmu.edu*
- *C. Xu is with the Center for Energy-Efficient Computing and Applications, School of EECS, Peking University, Beijing 100871, China. E-mail: chenren@pku.edu.cn.*
- *T. Vu is with the Department of Computer Science and Engineering, University of Colorado, Denver, CO 80204. E-mail: tam.vu@ucdenver.edu.*
- *M. Gruteser, R. Howard, Y. Zhang, and N. Mandayam are with the Wireless Information Networks Laboratory (WINLAB), Rutgers University, North Brunswick, NJ 08902. E-mail: {gruteser, reh, yyzhang, narayan}@winlab.rutgers.edu.*
- *W. Yuan is with the Google, Mountain View, CA 94043. E-mail: wenjia.yuan@gmail.com.*
- *K. Dana is with the Department of Electrical and Computer Engineering, Rutgers University, Piscataway, NJ 08854. E-mail: kdana@ece.rutgers.edu.*

Manuscript received 3 June 2015; revised 10 Nov. 2015; accepted 6 Jan. 2016. Date of publication 28 Jan. 2016; date of current version 31 Oct. 2016. For information on obtaining reprints of this article, please send e-mail to: reprints@ieee.org, and reference the Digital Object Identifier below. Digital Object Identifier no. 10.1109/TMC.2016.2522967

1536-1233 © 2016 IEEE. Personal use is permitted, but republication/redistribution requires IEEE permission.
See http://www.ieee.org/publications_standards/publications/rights/index.html for more information.

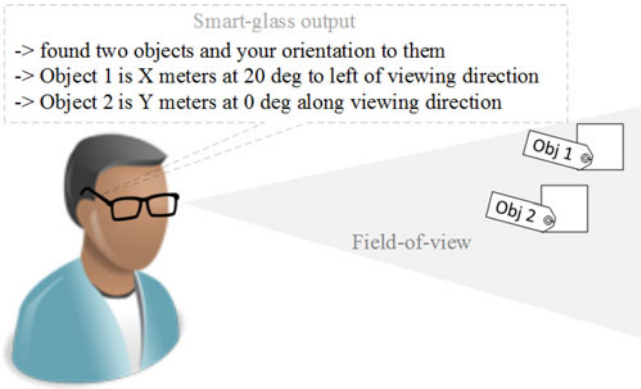

Fig. 1. Illustrating in-view recognition using smart-glasses.

analyze the camera footage from a mobile device to recognize objects, which works best with well-known landmarks [9] or previously recognized subjects/objects [10]. The accuracy of this approach degrades, however, as lighting conditions deteriorate, the number of candidate objects/subjects becomes very large, or the objects themselves look very similar (e.g., boxes in a warehouse). More importantly, camera operations are energy intensive and usually not optimized for long-term usage on battery power. Active approaches involve tagging objects of interest with a transmitting device, that emits signals carrying a unique identifying code pertaining to the object. An equivalent receiver that communicates with the transmitter recognizes the tagged object by decoding the identification code.

There has been a multitude of work in the active radio-frequency identification (RFID) community for positioning and localizing objects using radio signal strength (RSS); particularly indoors. However, using radio signal strength for positioning is very challenging due to multipath effect. Recent work on active RFIDs [11], [12] have been able to promise precise positioning accuracies. However, these techniques trade-off accuracy with energy-consumption and/or cost in consideration for smart-glass usage. For example, in [12] directional antennas provide the key benefit of accurately triangulating the reflected frequency-modulated continuous wave (FMCW) radio signal paths and pin-pointing the object's position. However, this approach may not be suitable for smart-glasses due to the high power consumption for generating the FMCW signals. In [11] though the solution achieves very precise positioning, the use of a R420 RFID reader makes the approach to be power consuming and expensive for integration on smart-glasses.

Apart from RF other possible modularities for positioning include using ultrasound. With its low propagation speed, ultrasound signals allow precise time-of-flight based angle-of-arrival (AoA) estimates. However, this is achieved at the cost of increased receiver size (5–10 cm) and a significant amount of energy to overcome it's exponential path-loss propagation; as opposed to inverse-square for electromagnetic waves. In addition, recent works on eye-gaze tracking systems [13], [14] have shown to precisely track what part of the scene is the user's eye concentrated upon. However, this solves only one dimension of the recognition problem as the identity of the object that the user is actually "looking at" still remains unknown. In general, solutions so-far in the eye-gazing have required intricate or expensive hardware design.

A hybrid radio-optical beaconing approach. We address the low-power challenge in-view recognition problem through a hybrid radio-optical system design. Our design adopts the active approach and integrates near-IR (infrared) based orientation tracking with object identification using active RFIDs. The key component of our design is the radio-optical signal, or *beacon* as we will refer to, which is an ensemble of a radio packet and an IR signal pulse. Infrared signals, due to their high directionality, can lead to precise orientation tracking through angle-of-arrival and distance estimation with a relatively small receiver, due to their small 850 nm wavelength. The main advantage of our approach is that it efficiently minimizes energy consumption by synchronizing the IR link between the transmitter and receiver using a RF side-channel. This enables the receiver to know exactly when to expect the IR pulse and thus allows for using extremely short IR pulses (order of μs) due to tight synchronization between the transmitter and receiver. The receiver uses the ratio of the received IR power over a photodiode array to determine the angle-of-arrival of the signal and uses the sum of absolute received powers to estimate the distance between the transmitter and receiver. The radio link is used to communicate the identity through a unique code embedded in the radio-packet.

Advantages of proposed approach. The key contribution through our proposed design in this paper is the *radio-synchronized IR signaling*. Since the transmitter and receiver are synchronized it allows to operate both, the IR and radio links at extremely low power thus optimizing battery usage. In addition to reduced energy consumption, short IR pulses also lead to a simplified IR receiver design – instead of requiring an infrared communication receiver (such as in TV remote controls), a synchronized energy detection circuit suffices. Existing IR technologies ([15], [16]) typically trade off energy consumption with range and/or beamwidth (angular-range). For example, GigaIR [16] can achieve low energy consumption with extremely narrow IR beams, but only within very short distances (i.e., tens of centimeters). The synchronization between the radio and optical links also makes it possible to estimate the 3D spatial position coordinates of the tags using a single radio-optical beacon sample at the receiver.

In summary, the key contributions from our work in this paper are as follows:

1. We propose a solution for low-power in-view recognition of objects in space using a hybrid, radio-optical, beaconing approach. We propose a design that leverages the high directionality of IR signals for precise orientation tracking and low-power nature of RFIDs to communicate identity. We develop a synchronization protocol to reliably associate the IR signals with the radio link and enabling use of extremely short IR pulses conserving significant battery power through synchronous transmission and reception.
2. We implement a proof-of-concept prototype radio-optical beaconing system. We design he hardware and software of a radio-optical transmitter (tag) and receiver, and develop a positioning application on an Android smart-glass heads-up display that integrates with the receiver. The application identifies

TABLE 1
Comparing Different Positioning Technologies

Technology	ID	AoA accuracy	Range	Size (order of)	Battery life
RF	ID from data-packets	low	NLOS(<100 m)	few cms	months-years
Ultrasound	require side channel	high	LOS (<14 m)	few inches	months
Camera	image recognition	high	LOS (10 s of m)	mm-10 s cm	1–2 days
IR	encode bits as pulses	high	LOS(<10 m)	mm-cm	few days

Size of cameras can trade off with speed and image quality.

and positions objects in user's view fit with the radio-optical tag.

3. We evaluate the tag and receiver power consumption. We estimate that our tag design can offer battery lifetimes upto few years. We discuss the effect of important design parameters such as IR pulse length, beaconing period and beaconing rate on tag power consumption. Overall, we show that the tag and receiver are least 1.5× power efficient compared to prior tag based approaches.
4. We experimentally evaluate the positioning accuracy of our approach in real-world settings with tags fit onto objects and the receiver fit on eye-glasses. We show that AoA and distance estimation errors are limited within 1–2 degrees and 40 cm respectively, over a maximum range of 9 m. We also show that our prototype in-view recognition application on the smart-glass receiver achieves 97 percent accuracy.

The outline of this paper is as follows: Section 2 motivates our work in this paper; Section 3 presents the system design in detail; Section 4 describes the prototype design and Section 5 discusses evaluations and results; Section 6 discusses challenges and limitations of this work and Section 7 concludes the paper.

2 Motivation

2.1 Requirements of In-View Recognition

We identify three key requirements for the task of in-view object recognition using smart-glasses.

Low battery power consumption. Due to the continuous operation requirement on battery powered devices it is critically important to significantly minimize power consumption on both, the tag and receiver. In this work we consider that both, tag and receiver, are battery operated. We envision that the tags will be fit onto mobile objects. The number of tags required for a finite space may be large and thus it becomes imperative to significantly reduce tag power consumption such that their batteries can operate for long durations (months or even years) without needing replacement. In this work we take a standpoint that the receiver will be fit onto a wearable device such as heads-up displays or eye-glasses. Though it is usually much easier to recharge the wearable devices than the tags it is fundamentally important that the in-view recognition modules on the receiver can operate low power usage considering it must be kept switched ON for long usage durations of hours or days.

Precise orientation estimation. The task of identifying multiple objects within a user's view demands precise orientation towards these objects (*knowing how far and at what orientation is the object in user's view*), and association between object spatial positions and their identity (*knowing what it is*). Determining such orientations is similar to estimating angle-of-arrival of a signal from the object to a reference point on user's view. Accurate object tracking will have very low AoA error tolerance. For example, with 1m spacing between objects and 3 m distance between the user and the objects, the AoA error tolerance so as to distinguish the objects in the user's view is about 10 degrees.

Small size. It is important that an enabling technology must make it possible for miniaturization; transmitters as integrate into small mobile objects and receiver as it must be integrated into a wearable device. For mounting a receiver onto smart-glasses, desirable receiver sizes will be of the order of centimeters and smaller.

2.2 Advantages and Limitations of Candidate Technologies for In-View Recognition

In Table 1, we compare individual candidate technologies for solving the in-view recognition problem. AoA using RF signal strength alone is very challenging due to the multipath nature of radio signals—the angle resolution is fundamentally limited to a few radians. However, communicating information through radio signals consumes very less power.

Ultrasound signals are good candidates for AoA estimation due to their low propagation speed; enabling precise ranging through accurate time-of-flight estimates. However, ultrasound transducers are costlier than radio antennas and ultrasound receivers have minimum size requirements due to its relatively long wavelength; the minimum distance of separation between ultrasound receivers on an array is from few cm to tens of cm.

The highly directional nature of optical signals makes it a viable candidate for accurate AoA estimation. Visible light based systems, e.g., using cameras, perform accurate pose-estimation to determine AoA based on preset markers, but cameras are energy intensive and unsuitable for long-duration operations.

IR signals, that are unobtrusive as they are invisible to the human eye, can give precise AoA estimation and ranging [17], [18]. However, IR wireless communication is much less energy efficient than RF because it has to overcome much higher ambient noise levels than in the RF spectrum. This means that transmitting even 1 bit of information using IR communication will incur significantly higher energy than RF.

2.3 Why Use Radio-Optical Approach

We learned that adopting a single technology will result in trading off accuracy with battery power or vice-versa. In this regard, we explore a strategy that blends multiple technologies leveraging the key advantage(s) of each

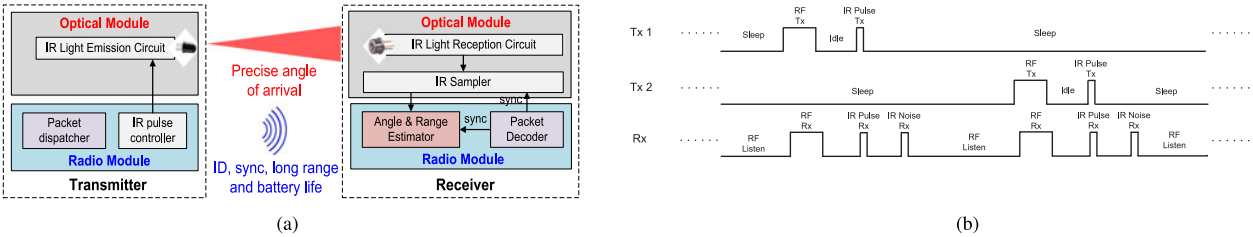

Fig. 2. (a) Radio-optical beaconing system architecture. The IR beacon is used for accurate positioning through AoA, while synchronization and ID communication is through radio. (b) Timing diagram of the *paired-beaconing* protocol over one duty-cycle period; this example uses two tags.

technology. In particular, we propose to bring together radio and optical (using IR signals) technologies to address the in-view recognition requirements. This hybrid approach can provide the following key benefits:

Energy efficiency through radio synchronization. With the availability of two orthogonal technologies it will be possible to use one to synchronize communication on one link using the other. Our proposed approach uses the radio link to synchronize the IR link. This approach enables to use extremely short IR pulses for positioning and the radio link for communication, conserving considerable battery power.

Precise AoA using IR signal strength. IR signals due to their high directionality, can provide robust AoA estimation and ranging. It is possible to determine AoA and distance using the signal strength alone of an IR pulse of a predetermined duration (kindly refer to the Appendix, which can be found on the Computer Society Digital Library at http://doi.ieeecomputersociety.org/10.1109/TMC.2016.2522967), for derivation of AoA estimation using IR signal strength). IR signals are unobtrusive for humans and do not travel through obstructions, reducing the likelihood of including objects that are not directly within sight.

Size reduction through simple design. The small size of RFIDs and the simplicity of the required IR circuitry (a pulse generator at the transmitter, energy detector at the receiver) makes the design conducive for miniaturization.

2.4 Applications

The radio-optical approach for the in-view recognition problem can enable plethora of applications requiring continuous battery usage and precise recognition. In particular, when integrated with smart-glasses, users can be navigated to tagged objects in stores, books in a library, items in a grocery store or warehouse etc. Smart-glasses can also be used to distinguish and locate different objects in a cluttered environment. In addition, precise orientation mapping becomes critical for augmented reality applications where physical world 3D coordinates of objects need to be mapped to the user's view. Apart from standalone smart-glass applications, the radio-synchronized IR signaling approach can be used as low-power signaling technique, for example in RFID systems or light based communication systems.

3 Radio-Optical Beaconing System Design

In-view recognition of an object on smart-glass using our radio-optical beaconing system will involve tagging objects of interest with a radio-optical tag that communicates with an equivalent radio-optical receiver fit on the smart-glass. The radio-optical receiver identifies the radio-optical tag using the unique identification code transmitted through a radio link. The receiver determines it's spatial orientation with the tag by estimating the AoA and distance using the signal strength of the IR signal received through an optical link. By associating the orientation with the identity of the tag, the receiver positions the tagged within the user's field-of-view.

The key novelty of our proposed design is the *radio-synchronized IR signaling* approach. In this approach, the radio and the IR links between the tag and the receiver are synchronized through our proposed synchronization protocol. This protocol enables the receiver to know exactly when the IR signal will be transmitted from the tag. This timing information is computed at the receiver by detecting the arrival time of the corresponding radio packet. Through this radio-synchronized IR signaling approach we mainly leverage three key benefits to address the in-view recognition problem for smart-glasses:

- *Extremely short IR pulsing:* It is possible to use an extremely short IR pulse for IR signaling hence providing significant gains in energy efficiency,
- *Synchronized radio and IR links:* It is possible to estimate precisely the arrival time of the IR signal using an orthogonal (radio) side-channel enabling easy association of orientation with identity and also avoiding a dedicated synchronization circuitry, and
- *Collision avoidance:* It is possible to preserve the accuracy of orientation estimation even in multi-tag environments by avoiding IR signal collisions using precise IR signal arrival time estimates through radio-synchronized link.

We illustrate our proposed radio-optical beaconing system through the architecture diagram in Fig. 2a. We describe the IR link through our model illustrated in Fig. 3. In this model, we consider a three-element photodetector IR receiver that samples IR signals from an IR LED transmitter on the radio-optical tag. We define AoA as the angle between the receiver surface normal and the vector connecting the transmitter and the depicted reference point at the center of the photodetector array; on the horizontal plane (azimuthal) as θ and vertical plane (polar) as ϕ. The photodiodes are rotated by an angle δ from the surface normal such that the angle between the LED and the vector in direction of photodiodes 1, 2 and 3 is $\theta - \delta$ and $\theta + \delta$, and $\phi + \delta$, respectively. We use δ_{ir} to represent the width of the IR pulse and d as the distance along the viewing axis between the tag and receiver. We refer the reader to the Appendix, available in the online supplementary material, for the mathematical derivation for determining AoA (θ and ϕ) and distance (d) from IR signal strength.

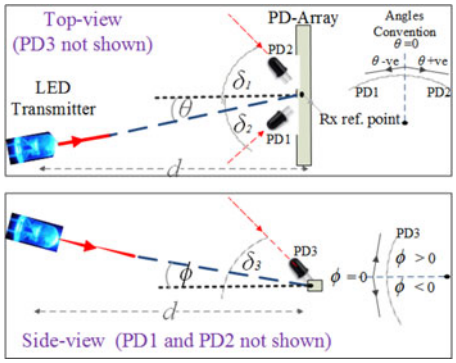

Fig. 3. Diagram to illustrate the model of infrared link in our design. A single LED transmitter – three element PD (photodetector) array receiver model ($\delta_1 = \delta_2 = \delta_3 = \delta$).

We will now describe the key aspects of our design in more detail.

3.1 Extremely Short IR Pulsing

We minimize the transmission period of the IR signal to the point where it can no longer be used for communication purposes but is still detectable for estimating AoA and distance. Theoretically, a single short IR pulse with maximum peak power, like an optical strobe light, can be detected even at 9 m and at very low *average* energy consumption due to its extremely short duration. The challenge, however, lies for the receiver in detecting *when* such a signal was transmitted. Adding a preamble, such as in communication systems, for detecting the signal would require multiple and possibly longer IR pulses, which leads to higher energy consumption. High (energy) efficiency detection therefore depends on proper timing—by enabling the detector only when the optical pulse is expected, we can eliminate much of the effect of background noise.

3.2 Synchronizing Radio and IR Links

We address the timing problem through a protocol that synchronizes the corresponding IR and RF signals at the tag and receiver. For simplicity, we will refer to this protocol as *paired-beaconing*. As illustrated in the timing diagram for *paired-beaconing* in Fig. 2b, the radio-optical tag periodically transmits a RF packet. Following the transmission of a RF packet, after a very short predetermined time-interval (known to the transmitter and receiver), an IR pulse is transmitted. The receiver uses the end of the radio packet as a reference to synchronize with the incoming IR pulse, and then samples the received signal from the IR signal receptors (photodiodes) over the expected pulse duration. It also takes an additional noise measurement after the pulse duration to calibrate for ambient noise-floor. The IR pulse itself carries no bits of information—the tag identity information is included in the preceding radio packet.

3.3 Collision Avoidance

A typical use-case scenario for our system would involve multiple tags transmitting to multiple smart-glass receivers. This implies that it is possible that the radio-optical beacons may collide in time as the radio links use the same bandwidth. Tags may also be placed very close to each other in space, and depending on the radiation pattern of the IR LEDs used the IR signals from multiple tags may collide in space. In such collision scenarios interference between tag transmissions will lead to erroneous recognition results. In our design we use the radio-optical approach to avoid such temporal as well as spatial collisions.

(a) Temporal collisions: The paired-beaconing synchronization mechanism ensures that each IR pulse is tightly associated with its corresponding radio packet carrying the identity. We limit the time-interval between the end of radio packet transmission and start of IR pulse to be extremely short (order of few micro-seconds) and duty-cycle the radio-optical beacon transmissions periodically. We keep the tag transmissions to be independent across each duty-cycle of a single tag and across transmissions from multiple tags. In that case, based on the model from Firner et al. [19], the probability of two tag transmissions colliding in our system is equal to the probability of the radio packets colliding within the radio packet duration in each duty cycle. For example, the collision probability our prototype system that uses 500 μs long radio packets duty-cycled at 1 sec intervals will be less than 5 percent for 100 tags. If the density of tags is increased, to retain such low collision probabilities, the radio packet duration must be decreased or the duty cycling interval must be increased.

(b) Spatial collisions: Radio packet collisions in space are implicitly avoided by the use of a correlation receiver on the receiver radio module. No specific measures are required to handle spatial interference of the radio signals. However, since the IR receiver uses simple energy detection, concurrent IR radiations from different tags can result in interference if the tags are spaced very close to each other. In our system, since tag transmissions are independent and that the radio transmission and IR pulsing are tightly synchronized, collisions of the IR signals in space can only occur if any two radio packets collide in time. By synchronizing the IR signals with the radio link the spatial collision problem can be virtually analyzed by studying the temporal collisions. For example, if two tags were placed next to each other with sufficient spatial separation, the tags can still be distinguished at a receiver as long as the radio packets from the two tags do not collide in time. The probability of such collisions will be very small considering the independent tag transmissions with extremely short packet durations and long duty-cycle intervals. However, the accuracy of positioning will impact the minimum spatial separation required between two tags in space. As we evaluate in Section 5, this minimum separation is about 20-40 cm for our prototype system.

Potential advanced protocols. In this paper we operate the smart-glass receiver in a passive always-ON mode. In case of a collision, the receiver discards the packet and waits for the next duty-cycle. Collision analysis with multiple smart-glass receivers in our system would essentially be similar single receiver use-case since the smart-glass is not enabled for transmission and that reception of each tag transmission is independent and identically distributed. Enabling the smart-glass units to transmit, by using a transceiver on the smart-glass, the communication between the tags and the smart-glasses can be further enhanced for improved collision avoidance through feedback protocols. On the transmitter end, in reality, the number of tags within the context of the smart-glass user's foveal view [20] may be limited. For

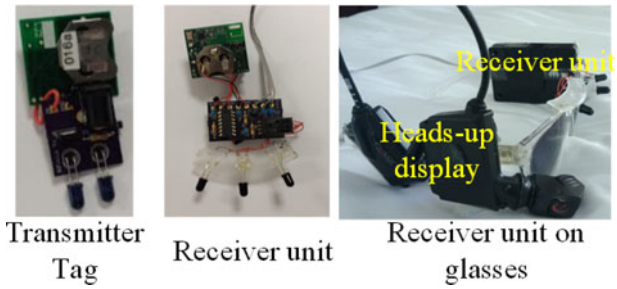

Fig. 4. Prototype tag and receiver (batteries not shown). The tag is 4 cm in largest dimension. Receiver unit is sized ($l \times w \times h$) at 5 cm $\times$ 4 cm $\times$ 3 cm.

example, there may be 1,000 tags in a storehouse but the number of tags within the context of the user's view may be limited to those within one or two racks. This means that it may be possible to adopt intelligent filtering schemes to limit the number of tags the receiver would decode from in a particular space and time context. For example, one simple approach may be to engineer the system to filter the tags not within a calibrated radio range (RSSI is less than a threshold).

In this paper, we have focused on emphasizing the benefits of radio-synchronized IR signaling, leaving advanced hardware and protocol design challenges for future work.

4 Prototype Design

We have prototyped the radio-optical tag, and a wearable receiver unit as shown in Fig. 4. We mounted the receiver unit on eye-glass together with a RECON Instruments MOD LIVE heads-up-display that runs Android. We developed a positioning application on this prototype receiver that uses our recognition framework. A proof-of-concept demonstration video of the positioning application [21] using our prototype can be viewed at http://www.andrew.cmu.edu/user/ashwina/papers/bifocusvid.mp4. In this demo we measured the total response time of the app to be 25 ms; we define response time of the app as the time duration between the instance the tag transmits a radio-optical beacon and the instance the Android heads-up displays outputs the identity and position) on its screen.

4.1 Radio-Optical Tags

The transmitter tag consists of a RFID module that is used for the radio communication as well as triggering the pulse input to an IR LED. To be detectable at maximum distance the LED has to be operated for maximum light emission. The LED achieves maximum light emission when the current (voltage) across the LED is 1A (2.5 V). As shown in the tag circuit diagram in Fig. 5, a MOSFET amplifier and an appropriate series resistor ensured that the current across each LED was maintained at 1 A. To maximize range, we used two near-IR LEDs [22] on the prototype tag. A high-energy pulsed LED emission requires a large spike in energy which cannot be achieved if powered by the same power supply of the radio. So we use an independent 9 V battery supply for the driving the LEDs and use a capacitor to prevent a sudden large voltage drop when the LEDs are activated. The 9 V power supply can be avoided by using a lower voltage battery along with a voltage step-up circuit. Over-driving the LED for maximum range can also be avoided by using multiple high power LEDs at nominal current drive. We reserve such design considerations for future.

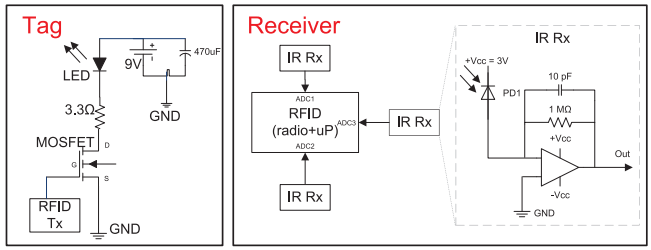

Fig. 5. Circuit diagrams of prototype tag and receiver.

The RFID module on the tag contains a CC1100 radio and a MSP430 microprocessor and powered by a CR2032 – 3V lithium coin cell battery. The radio operates at a data rate of 250 kbps with MSK modulation and a programmed RF output power of 0 dbm. In each duty cycle, the radio broadcasts a 12-byte packet (4 bytes of preamble, 4 bytes of sync, 1 byte of packet length, 3 bytes of tag id + parity bits), waits for short delay (measured to be at least 500 μs: over-the-air packet time of 380 and 120 μs hardware delay), triggers a 3 V pulse for a duration of $\delta_{ir} = 10$ μs on one general purpose I/O pin connected to the MOSFET gate, and goes back to its sleep mode. The radio wakes up every $\tau = 1$ sec and repeats the transmission.

4.2 Radio-Optical Receiver

The front end of the receiver consists of three Silicon photodiodes [23]. Two of them are horizontally spaced by 3 cm and mounted with 40 degrees separation (at half-power angle $\delta = 20$ degrees, symmetrical on each side); the third is placed 20 degrees off (on top) the horizontal plane formed by the other two. With this setting, our receiver achieves an angular-coverage of ± 20 degrees, and can be increased by placing more photodiodes in the receiver array. To amplify the detected IR signal on the photodiodes, we use an opamp (operational amplifier) and choose the resistor and capacitor values in the opamp circuit such that the rise-time (proportional to the time-constant – the product of resistance and capacitance across the opamp) is much less than the IR pulse length, so as to ensure maximum IR light energy accumulation over the pulse detection period at the receiver.

The receiver RFID module contains a CC1100 radio and a MSP430 microprocessor, similar to the radio-optical tag. Each photodiode's analog output from the opamp is wired to each of the three 12-bit analog-to-digital converter (ADC) input pins of the microprocessor. We power the radio using one of the 3 V supplies to the opamp (the opamp requires a +Vcc and $-$Vcc supply). We programmed the radio to stay in always-active receive mode ready for receiving the radio packets and IR beacon. Upon a successful packet reception the signal from the photodiodes are sampled at each ADC, and at a time instance after the end of packet reception – subject to a small hardware delay. The ADC sampling duration is set equal to the length of the IR pulse. The receiver identifies each tag through the unique transmit ID encoded in the radio packet.

The sampled ADC voltage readings correspond to the received IR signals; let us denote them as V_{h1}, V_{h2} and V_v.

TABLE 2
Energy Consumption and Average Power of Prototype Tag for a 10 μs IR Beacon (Two LEDs on Tag) and Radio Transmitting a 12 Byte Packet at 250 KBPS Every 1sec at 1 mW (0dbm) Output Power

State	Duration δ [μs]	I_{bat} [mA]	Energy [μJ]
idle	800	2.95	7.08
RF transmit	500	15.52	23.28
IR transmit	10	527.8	48.029
sleep	998,688	0.0007	2.097
Total energy	E_{tot}		80.486
Avg. power	$P_{avg} =$	E_{tot}/τ	80.486 μW

Energy $= V_{bat} I_{bat} \delta$, *where* $V_{bat} = 3V$ *for radio module and 8.1V for IR,* τ *=1 s; these also include the microprocessor's consumption since we account for the total battery drain.*

After obtaining the signal samples, the background noise (voltage) is measured by sampling the photodiode outputs after a 60 μs delay (10 μs of opamp delay plus 50 μs pulse fall-time), and for a duration equal to length of IR pulse. Let us denote the noise readings as N_{h1}, N_{h2} and N_v. Since the load resistance is the same for all the voltage readings, the angle and distance are estimated by substituting the numerical values of $V_{h1} - N_{h1}$, $V_{h2} - N_{h2}$ and $V_v - N_v$ values into I_{h1}, I_{h2} and I_v respectively (refer to equations in Appendix, available in the online supplementary material.

5 Experimental Evaluation

We conducted extensive experiments in a well-lit academic laboratory environment using our prototype tags and eyeglasses fit with the receiver in different real-world use-case settings. We evaluated the performance of our system based on the following metrics:

(1) Power consumption: We evaluate the battery power consumption of the tag and receiver units separately. We also conduct micro-benchmark evaluations to study the effect of different parameter choices on battery power consumption.

(2) Recognition Accuracy: We evaluate accuracy of our recognition framework at two levels:

(a) Orientation estimation accuracy. We evaluate the accuracy of orientation estimation through AoA and distance estimates. We use the AoA and distance estimation error metrics to represent accuracy of orientation estimation.

(b) Recognition accuracy of application. We conduct a macro-benchmark evaluation to study the end-to-end accuracy of the radio-optical in-view recognition system through our prototype smart-glass application.

5.1 Power Consumption

5.1.1 Tag Power Consumption

We compute the tag average power consumption as $P_{avg} = \frac{E_{tot}}{\tau}$, where τ is the beaconing period (duty-cycle duration). The total energy consumption E_{tot} of the tag is the cumulative amount of energy consumed by the three modules: microprocessor, radio, and IR. In Table 2, we report the battery energy consumption in different states of operation during a 1 s beaconing period. We measured the current draw from the battery source in different states of operation; separately for the radio and IR modules as they are powered by independent battery sources. We compute I_{bat} as total current in each state of the tag by integrating corresponding regions of oscilloscope readings from Fig. 6a for radio, and (b) for IR module. The 'idle' state in Table 2 includes the transitioning periods from 'sleep'to ON and vice-versa. Finally, we compute the tag average power consumption as $P_{avg} = \frac{E_{tot}}{\tau}$, to be 80.486 μW for a 1 second beaconing period in our prototype.

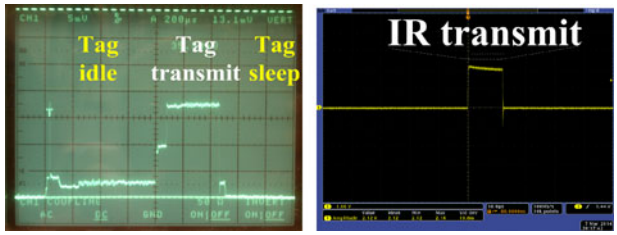

(a) (b)

Fig. 6. (a) Tag's radio module battery drain (voltage reading is across a 1 Ω resistor on an analog oscilloscope). (b) Tag's IR module battery drain (voltage reading is across a 3.9 Ω resistor on a digital oscilloscope.

Comparison with other prototypes. We compare the tag power consumption with two technologies; an IR remote control, and ultrasound [26]. For the remote-control, we determined the IR pulse period to be 10 μs and peak current draw at 50 mA from a 3 V (two alkaline AAA batteries) supply. We interpolate the effective pulse-period to be 1.04 ms for transmitting 13 bytes (includes preamble and ID), yielding a total energy consumption of approximately 150 mW. IR remote control technology is a less energy-efficient option for continuous operation in recognition applications. This is because the IR transmission will have to communicate a packet of bits where each bit corresponds to one IR pulse, thus keeping the battery on and draining the peak power for a longer duration. We eliminate the need for transmitting multiple IR pulses by using the radio channel to communicate the ID through a RFID packet. As can be seen from Table 2, the RF transmission to send an entire packet consumes less than half of the battery energy compared to transmitting a single IR pulse. In comparison with the ultrasound based positioning system implementation in SpiderBat [26], our transmitter achieves about 1.5x higher energy efficiency. SpiderBat integrates two orthogonal technologies (ultrasound and radio), however, the radio channel is only used as a side-channel to communicate ID but not for synchronizing the ultrasound transmissions.

5.1.2 Micro-Benchmarks for Tag Power Consumption

Tag Battery Life. Our power measurements indicate that the radio module and IR module consume about 33 μW and 49 μW respectively, for 10 μs IR pulse and 1 sec duty-cycling. Theoretically, this means that the tags can continuously operate for a lifetime of about 9 years on a 9 V alkaline battery (520 mAhrs capacity) beaconing once per second. In practice, the irregularities and limitations in the circuit (components, wiring etc.) and battery design can reduce the lifetime. In reality, we feel that with such low battery power consumptions, achieving a lifetime of at least few years is possible with the right choice of circuit and battery design.

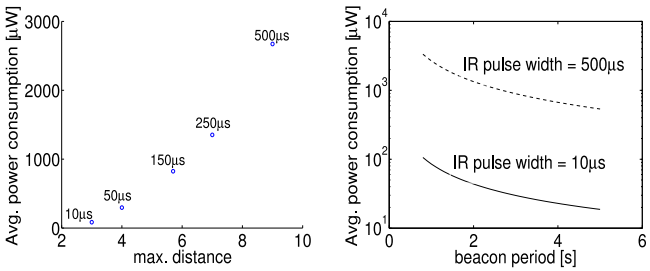

(a) (b)

Fig. 7. (a) Tag power consumption versus maximum distance of operation (range of the system) and (b) Tag power consumption versus beaconing period, for different IR pulse width choices.

IR Pulse Length. A large IR pulse length in our system implies more radiated IR energy. Increasing IR pulse length enables to increase the angular and/or distance range, but trades-off with the increase in battery power consumption. In Fig. 7a we plot the measured power consumption versus the distance range (maximum distance of recognition) in our system, for different IR pulse duration choices. As can be seen from Fig. 7a, a 10 μs pulse can achieve only a 3 m range. However, we believe that this range still useful for many short distance interactive applications on the smart-glasses; for example, interacting with posters in a conference or artifacts in a museum, searching for items in a shelf in a store.

Transmitter Beaconing Period. In Fig. 7b we plot the transmitter power consumption for different beaconing periods τ. The plot indicates that, for a 10 μs pulse, increasing the beaconing period (less beacons per unit time) is increased to 5 seconds considerably saves battery power. However, power savings is less pronounced when increasing the IR pulse length to 500 μs. We learn from Fig. 7b that the battery power usage monotonically increases with increase in number of beacons transmitted per unit time.

5.1.3 Receiver Power Consumption

The power consumption at the receiver includes that of the radio receiver and IR module. Table 3 compares the average power consumption of our receiver prototype with other existing technologies for positioning; ultrasound (US) and camera. Overall, we observe that our design can help achieve 2.5× more efficient on battery consumption on the receiver compared to other approaches using competitive technologies. Based on our measurements, the battery life (of a 3 V alkaline AA battery) for only the receiver operation is little less than 2 days. This ensures that it is possible to operate the smart-glass receiver in an always-ON mode for about two days without powering OFF.

TABLE 3
Comparison of Receiver Average-Power Consumption (P_r) with Other Possible Technologies [($*$ Uses Image Recognition, Subject to the Tagged Objects not Similar Looking, and Will Require at Least Two Image Frames to Avoid Aliasing), ($**$ Per-Image)]

Technology	Deliverables	P_r [mW]	Total P_r[mW]
Ours	AoA (θ, Φ) (IR)	9	
	ID+sync (radio)	81	90
Ultrasound [26]	AoA (θ) (US)	140	
	ID (radio)	100	240
Camera [27]	AoA(θ, Φ),ID*(image)	202.2**	202.2

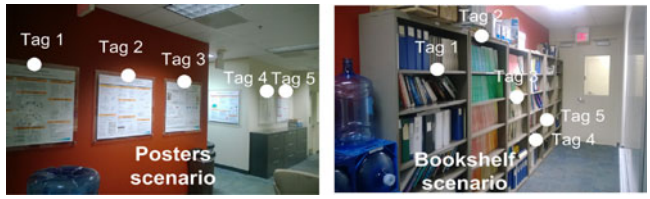

(a) Tags on posters (b) Tags on shelves

Fig. 8. Experimented application scenarios (a) *Poster* and (b) *Bookshelf*.

5.1.4 Discussion

Battery power usage optimization is key to any mobile system. In this paper, we have focused primarily on optimizing the battery power consumption of the tags, as we envision an environment where the tags are attached to mobile objects. We consider that the primary power source for a tag is from a battery on the mobile object and that it may not be possible to periodically recharge the battery on such objects. For the receiver, we take an optimistic approach and believe that a wearable device, such as the smart-glass, can typically be switched-ON on a need-to-use basis and can be recharged periodically. We believe that through strategical techniques such as conserved usage of tag and/or receiver for power savings is possible to further reduce power consumption. In this paper, we primarily emphasize the possibility of significant power savings for long-duration operations through the radio-synchronized IR signaling, enabling to operate the tags and receiver at much lower battery drain than other strong candidate approaches such as vision.

5.2 Recognition Accuracy

We conduct experiments by emulating a smart-glass user's behavior in four real-world application scenarios (as shown in Figs. 8 and 9):

(i). *Poster* (Fig. 8a), that represents a scenario where smart-glass users interact with advertisements or posters, (ii). *Bookshelf* (Fig. 8b), that represents a scenario where smart-glass users desire to locate a certain object such as shelf in a library or warehouse, (iii). *Office-Room* (Fig. 9a), that represents a scenario where smart-glass users desire to locate an object in a relatively large and neat office room where tagged objects are spread out, (iv). *Cubicle* (Fig. 9b), that represents a scenario where smart-glass users desire to locate items in a cluttered, small space, such as a cubicle or a medicine cabinet.

Experiment setup. To facilitate ground-truth angle measurements, we attached a camera recording video frames at 30 fps, fit with an IR lens (will refer to as IR-camera) onto the glasses as shown in Fig. 10. We use the camera for ground-truth AoA measurements; angle subtended by the light ray with the camera reference axis can be determined accurately from the pixel image coordinate of the imaged light emitter (captured as a white blob by the IR-camera) using camera projection theory [24]. We fit the photo-diode array onto the camera such that the reference axis of the photodiode array and camera are the same. This setup avoids errors due to any

(a) Tags on objects (b) Tags on objects

Fig. 9. Experimented application scenarios (a) *Office-room* and (b) *Cubicle*.

discrepancy in ground-truth measurements and movement of the array. For manual visual verification we also fit a smart-phone camera onto a helmet that was worn by the experimenter during the course of experiments.

5.2.1 Experiment Methodology

In each experiment scenario (*Poster, Bookshelf, Office-room, and Cubicle*) we used a total of five tags, that beaconed an IR pulse of width 10 μs every 1 second – a 10 μs IR signal integrated over the 33 ms frame period (30 fps) was detectable by the CMOS sensor of the camera, due to high light energy output from the LED. All the data, along with timestamps, was collected on a linux laptop with the camera connected through USB. We collected a total of 15,000 data samples (over 4 hours of experimentation), where the experimenter (one of the authors), wore the prototype glasses, and performed the following actions in each scenario:

(i) *Poster*: Read a poster, from a distance of 2 m, for a few minutes and move to the next. Before moving to the next poster, turn head to look at the subsequent poster from the current location and then walk to it.

(ii) *Bookshelf*: Search to locate a particular bookshelf. Here, first try to locate the shelf (standing 1.5 m away from the shelf and looking up or down) and then make lateral head-movements to emulate searching for a particular item on that shelf. Repeat the same exercise for the subsequent tagged shelves.

(iii) *Office-room*: Search for a particular tagged object in the room, gaze at it for a few seconds. Repeat the same for other tagged objects. During the course of the experiments, the experimenter is seated on a chair 1.5 m away along the 0 deg axis facing Tag 2 in Fig. 9a.

(iv) *Cubicle*: The actions in this experiment are the same as in the *Office-room* scenario, but with the tags placed in a more cluttered space. During the course of the experiment, the experimenter stands 1.5 m away along the 0 deg axis of Tag 3 in Fig. 9b.

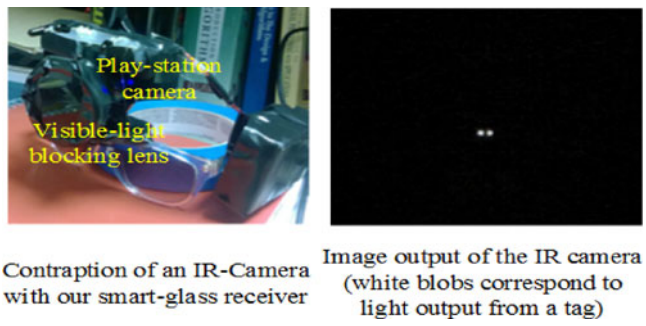

Fig. 10. The image on left shows the contraption we used for our experiments. We fit an IR lens onto a SONY play-station camera and fit onto our smart-glass prototype. The image on the right shows the output of the camera fit with the IR lens.

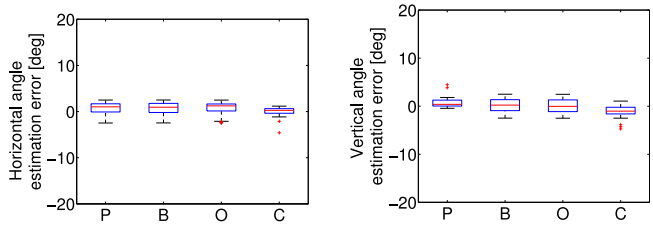

(a) Horizontal (b) Vertical

Fig. 11. Angle-of-arrival estimation error for the four application scenarios (P, B, O, and C refer to *Posters*, *Bookshelf*, *Office-room*, and *Cubicle* scenarios, respectively).

Why we chose these scenarios? The key reason behind adopting these scenarios is that each of these cases differ primarily in the arrangement of the tags. Different tag arrangements can lead to variety of 3D points in space for each tag and user head position combination. By letting the user make head movements while targeting at a tagged object, over a period of time, we populate a diverse and large number of 3D points which serve as the sample space for evaluating the accuracy in our system. The additional reason is that we also want to ensure that in such diverse environments, where the number of obstructions (possibilities of radio multipath) vary, our radio-optical beaconing approach can still provide accurate AoA and distance estimates for positioning. The *Poster* scenario emulates searching for tagged items that are arranged in an uniform order and are at almost the same distance from the smart-glass user when the user is directly directly looking at the tag (at 0 degree angle). The *Bookshelf* scenario also has an uniform order but the tags are at a different height from each other. The *Office-room* and *Cubicle* scenario represent a more disordered arrangement of tags. However, in the *Cubicle* scenario the signals may encounter more obstructions and thus possibility of IR reflections are higher resulting in positioning errors.

5.2.2 Results for Orientation Estimation Accuracy

Angle-of-arrival: In Figs. 11a and b, we plot the errors in horizontal and vertical angle estimates, respectively. By analyzing the cumulative distribution of the angle error data we found that the median error is 1.2 degree and 80 percent of the errors are contained within 1.5 degrees, and maximum error is at 2.2 degrees. We also observe that these error distributions are consistent for both horizontal and vertical angles. We believe that an AoA error of 1.2 degree is acceptable for most applications. We note that the angle estimation errors reported here also include the deviations in the ground truth angle measurements due to head movement. We examine this further in Section 6.

Distance: We first evaluate the distance estimation accuracy in each spatial dimension through a controlled experiment where the experimenter, wearing the glasses receiver, positioned the head so as to look at only one tag and did not make any head movements. Two sets of data were collected, where in each, one angular dimension (horizontal or vertical) was fixed (to 0 deg) and the other changed; the perpendicular distance between the experimenter and tag was fixed at 3 m. We report the distance error estimates from

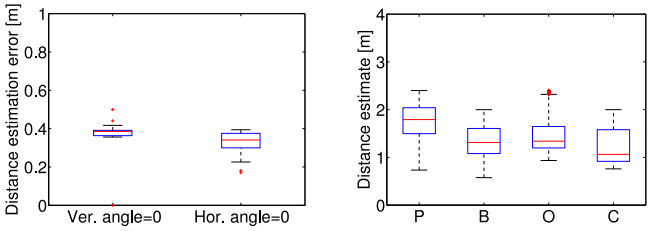

(a) (b)

Fig. 12. (a) Distance estimation error for head-worn receiver setup with no head movements. (b) Distance estimates for head-worn receiver setup in four application scenarios with free head movements (P, B, O, and C refer to *Posters*, *Bookshelf*, *Office-room*, and *Cubicle* scenarios, respectively).

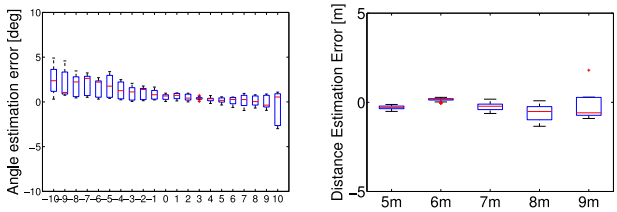

(a) Angle estimation error (deg) vs. Horizontal angle (deg) (b) Distance estimation error (m) vs. Distance (m)

Fig. 13. (a) Angle estimation error for calibrated (fixed) setup. (b) Distance estimation error for calibrated (fixed) setup.

this controlled experiment in Fig. 12a and observe that the median distance error is within 40 cm in both cases. In Fig. 12b we plot the distance estimates (instead of distance estimation errors) from the experiments in the four application scenarios. We observe from Fig. 12b that the distance estimates are in agreement with the actual distance the experimenter maintained during the course of experiments (2 m for the poster scenario and 1.5 m for others). We also observe that the median errors are within the 40 cm value as evaluated through the calibrated set-up.

AoA and Distance from calibrated experiments: We have also conducted experiments to validate our AoA and distance estimation accuracy in a tightly calibrated setting. In this set of experiments, we marked locations on the laboratory floor for ground-truth angle and distance measurements. The measurements spanned -10 to 10 degree in 1 degree spacings on the horizontal ($\theta_{ver} = 0$) and from 5 to 9 m in 1 m steps. At each marked test points, we positioned the tag and collected 60 consecutive beacon samples. We then repeated the entire procedure 5 times, yielding a total of 300 samples per test point. We performed our evaluations for the tag beaconing period of 1 sec and an IR pulse length of 500 μs to maximize range. In this experiment, the receiver glasses and transmitter were both positioned (fixed) on a crate at an equal height of 60 cm from the floor. As can be seen from the angle error and the distance error plots in Figs. 13a, b respectively, the median horizontal angle error (of 1.2 degrees) and median distance error (of 40 cm) from the application scenario experiment, are in good agreement with results from the calibrated setup.

5.2.3 Results for Application-Level Recognition Accuracy

We define the recognition accuracy at the application level as the ratio of the total number of successful recognition events over the total number of events, collective of all the trials. We associate three key parameters with each tagged object: a unique object ID, AoA, and distance between tag the receiver. For our evaluation, if the receiver decodes the ID, and if angle and distance estimation errors are contained within 2.2 degree and 40 cm respectively, we considered it as a successful recognition event or a true-positive. We select these numbers based on our empirical observations from the results obtained for AoA and distance estimation. Our results indicate that, on an average, a tag within a user's view is successfully recognized 97.5 percent of the time. The individual true-positive rates for the scenarios, *Poster*, *Bookshelf*, *Office*, *Cubicle*, are 97.1, 98.025, 97.5, and 98 percent, respectively. We feel that such recognition accuracies are comparable with sophisticated vision based approaches available today [25].

We also observed recognition failures. Failure in recognizing a tag are triggered by false-positive and false-negative events. False-positive events mainly occur due to reflections of the IR signal. We observed that, across all four scenarios, the false-positive rate was within 2 percent, among which the *Poster* scenario had the most false-positives due to reflections of the IR signal from the smooth plexiglas surface. We refer the reader to Section 6 for our experimental findings that characterize the IR reflection level from various common surfaces. On the other hand, false-negatives occur when the RF or IR beacon is lost. We observe a false-negative rate of 0.5 percent across all the scenarios. We note that this rate may go up when the tag density increases, but we expect efficient medium-access protocols in future designs can effectively minimize RF collisions under a reasonable density. Since the total number of true events far exceeded the false events in the collective set we observed high precision and recall values of 96.5 and 98.4 percent respectively.

5.2.4 Accuracy versus Power Consumption Tradeoff

Intuitively, we expect that the accuracy of system will typically trade-off with power-consumption. To study and verify if such a tradeoff exists, we evaluate the application recognition accuracy and average power consumption using measurements from our prototype radio-optical in-view recognition system at different transmission rates. As can be seen from Table 4, the average power consumption monotonically increases with the transmission rate. This is not surprising because we do expect that the total number of radio-optical beacons transmitted per unit time will increase with the rate and thus drawing battery power at higher rate. We also observe that recognition accuracy increases with the transmission rate. This is because, as more samples are available at the receiver the errors are averaged out and accuracy improves. Hence, the trade-off between accuracy and power-consumption is thus evident from Table 4.

The main reason for such high accuracy in our system is due to fact that the probability of collision between beacons is minimized by the tight synchronization between the radio and IR links. However, mobility can induce errors into the system. In this regard, we make the following inferences

TABLE 4
Recognition Accuracy and Power Consumption for Different Beaconing Rates

Beaconing rate [beacons/sec]	0.5	1	2	5	10
Accuracy (S) [%]	98.5	99	99	99	99.1
Accuracy (M) [%]	95	97	98	98.5	99
Energy [μ W]	40	81	165	410	850

We consider the tag is static. Receiver may be static (S) or mobile (M). The IR pulse duration is 10 μs.

from Table 4 regarding effect of beaconing rate and receiver mobility on accuracy:

(a) Static use-case: In the static use-case where the receiver is placed on a table (fixed with no motion), the accuracy is almost constant across different beaconing rate choices. This is because, since both the tag and receiver are fixed, the number of erroneous and false events are very small. Therefore, even though more number of samples are obtained at the receiver the effective improvement in accuracy is minimal.

(b) Mobile use-case: In the mobile use-case, where the user wears the receiver on the head and makes head movements, it is possible that tag may not be in line of sight of the user all the time. Hence, the possibility of number of erroneous, false and missed events can increase. Providing more beacon samples by increasing the beaconing rate, enables the receiver to compensate for missed events and improve estimation errors.

5.2.5 Discussion

In our model, we relate to the 3D space using horizontal AoA, vertical AoA, and distance between the tag and receiver. These values represent the three independent dimensions of the object's position in space.

Our results indicate a 1.2 degree AoA error in horizontal and vertical dimensions. This translates to about 20 cm drift from the actual position in the horizontal and vertical coordinates, at 9 m distance. This means that the receiver can reliably distinguish the AoA from two tags if they are separated at least by 20 cm in horizontal and vertical dimensions. Since distance is in another orthogonal dimension, the ranging error of 40 cm cannot be directly translated to the minimum distance of separation between two tags on the horizontal and vertical planes. A 40 cm error in ranging will result in our system detecting the tag d m away along line-of-sight of the user to the tag within a $d \pm 0.4$ m accuracy. In our system, the ID, AoA and distance estimates are all obtained from one signal sample (one radio-optical beacon). The positioning accuracy can be further improved through improved noise reduction techniques for IR link and applying statistical learning on large receiver sample sets. Through improved collision avoidance, for example at the MAC layer, we can further reduce radio collision probabilities for large tag deployments.

Within the scope of this paper, we have presented a preliminary system design to achieve high positioning accuracy at low-battery power consumption. We have also shown the practical feasibility of an in-view recognition application with high positioning accuracies on a wearable device.

6 Limitations and Open Problems

Let us briefly discuss the limitations in our work, our explorations on trying to address them, and potential opportunities for future work.

Head movement. For head-mounted wearable receivers, our system only tracks head pose. The object a person is looking at, however, is also affected by eye movement. To understand the effect of head and eye movement, we sought to characterize how consistent head positions are when repeatedly looking at a series of objects. We fitted a laser pointer on the prototype glasses and one of the authors wore this contraption while repeatedly looking at objects on the wall. We recorded a video of the movement of the laser pointer. By analyzing this video footage and knowing the standing position and height of the user, we determined the effective angular deviations of the marker from the objects' exact position. We report the cumulative distribution of the data in Fig. 14a. Our experiment indicates a maximum of 1 degree and a median of 0.5 degree angular deviation between the head position and the object the person was looking at. We can infer that a head-mounted system would face this fundamental limit on angular accuracy—any higher precision would require additional eye tracking hardware.

Reflections. Due to the high energy on the IR pulse, reflections from smooth or shiny surfaces, can cause false detection of the beacons on the photodiode receiver. To study the signal strength of reflected beacons, we operated the prototype tag at 500 μs IR pulse width at different distances from three different reflecting surfaces: whiteboard, glass pane (see through), and dry-wall. The reflected signals were sampled by the photodiodes on the glasses receiver. The setup is illustrated in Fig. 14b. We conducted the experiment with no ambient lighting. Fig. 14c shows the graph of maximum

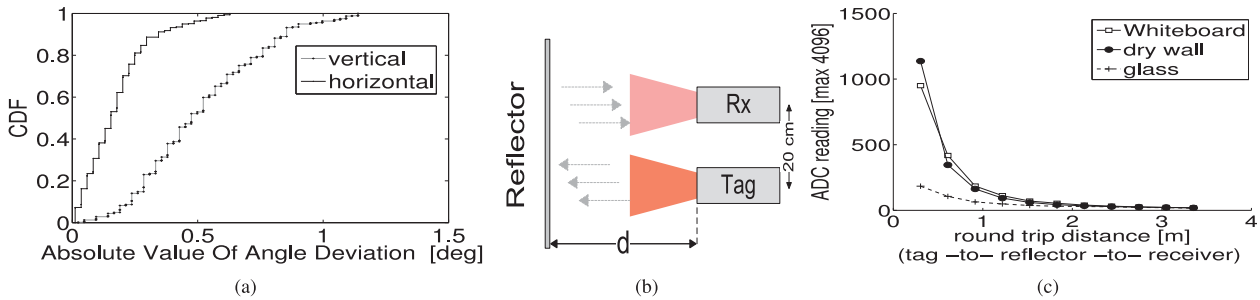

Fig. 14. (a) CDF of angle deviation for head-mounted receivers (experiment repeated for two users of same height). (b) Reflections experiment setup. We used distance of the tag (and Rx) from reflector surface, and accounted for the 20 cm spacing, in the round-trip distance computation. (c) IR signal strength from reflections versus round-trip distance.

of the three photodiodes' signal strength (as ADC readings) versus the total round-trip distance of the IR signal. Our measurements indicate that the effect of reflections (from typical indoor reflector surfaces) is negligible for round-trip distances greater than 3 m. Of course, for a 10 μs IR pulse this distance would be much smaller.

Object Switching Latency. The system may yield erroneous orientation estimates when the receiver suddenly switches between objects that it is tracking; for example when a user momentarily shifts head position. This happens due to the discrepancy of the IR signal strength on the photodiodes. Let $\Delta_t = t_1 - t_0$, where t_0 denotes the time instance when tag A goes out-of-view when the user starts to shift head position from tag A to tag B, and t_1 denotes the time instance when tag B is successfully located (within$\pm$2 degree error). We define the object switching latency in terms of the number of beacons as $(\Delta_t^{our-system} - \Delta_t^{camera-ground-truth})/\tau$. We determined the object switching latency of our prototype system through experimentation. Using a total of 100 head-turn events (switch head position from one tag to another) from our collective dataset, we determine that this latency is within a two beacon duration. This means that, when a user shifts head position it is possible to lock onto an object's signal within two extra beacons. The latency can be reduced by choosing a higher tag beaconing rate, but the trade-off is a larger battery drain (see Section 5.1).

Energy improvements. Reduction in power consumption can be achieved by reducing the pulse length and optimizing the circuit to eliminate noise sources. With this approach it should be possible to achieve larger ranges at about 10 μs pulse durations. Detecting optical pulses of 10 μs or less duration requires a high speed photo detector and high-speed photo integrator [28]. Here, the mechanical design requires replacing the complex mechanical layout with a carefully designed PCB and appropriate 3D shielding boxes on the board to avoid any electrical pick-ups. The 9 V battery in our current design could also be replaced with smaller batteries, such as three coin cells with simple circuit changes.

Size and monetary cost. The size of our prototype tag, is mainly governed by the size of the RFID tag used, and is 3.5 cm in the largest dimension. The size of the receiver depends on the placement of the photodiodes in the array, where the precision of placement and spacing can affect the accuracy of AoA estimates. We were able to achieve a 1 deg angular accuracy and 9 m range with a 3 cm spacing between each pair of photodiodes. Our prototype receiver unit is sized ($l \times w \times h$) at 5 cm $\times$ 4 cm $\times$ 3 cm. We believe that using surface-mount components on a printed circuit board (PCB) would reduce the size further. The hardware components (LED, photodiode, opamp, passive components) used in our design are all available off-the-shelf. Our design is not specific to the RFID used in our prototype. In essence, any commercial off-the-shelf RFID with access to their I/O pins will work. In our current prototype, the RFID cost is \$10 US while the IR circuitry cost only \$2 US. We believe that when considering large scale manufacturing the total cost (tag+ receiver) of the unit will be order of few US \$. Our work in this paper essentially presents a lab research prototype for addressing fundamental challenge of low-power in-view recognition. Integrating this design into commercial off-the-shelf devices will require more effort and may not be possible in an academic lab environment alone.

7 Related Work

RFID. A rich body of work exists in the area of positioning using RF-based ranging [29], [30]. However, it has been shown that radio signal strength based ranging is poor in high multipath environments [31] such as indoors. Kusy et al. [32] propose to address the multipath problem by using the doppler shift effects between a sensor and an anchor node, each transmitting at different frequencies. However, while the techniques improve the positioning performance, they require a large number of sensor nodes, making it infeasible in many practical settings. Wang et al. [33] have worked around the multipath problem in RF localization and achieved precise positioning by solving a space partitioning convex optimization problem. We differentiate our work from this prior work that our solution is very distributed as it does not require the tags to communicate with a centralized control or solver.

Ultrasound. The *Bat* [34] system is a well-known example which requires a large array of ultrasound (US) sensors. The *Cricket* [35] and *Calamari* [36] platforms integrate the US sensor into a RF wireless sensor platform. The downsides of ultrasound, however, are limited range (up to a few tens of feet), costly hardware and short battery life. For instance, ultrasound transducers are about an order costlier than the typical antennas used in RFID tags. *Medusa* [37], and *Spiderbat* [26] perform AoA based positioning with an array of US tranceivers, but are larger in size and less energy efficient than our prototype.

IR based positioning. Commonly used in robotics, where a robot is equipped with IR transceivers, IR is a well known candidate for positioning [17], [38], [39], [40]. However, due to the modulation, synchronization and demodulation circuitry involved, the power consumption of IR communication systems will be greater than a simple light beaconing. The *Active Badge* system [41], Firefly [42] are other examples of commercially deployed IR localization systems that use IR positioning, but are very energy intensive; making them infeasible for wearable systems.

Vision based systems. There are vision systems that are assisted by LED markers for robust identification [43], [44], however, such systems require extra processing for communication as simple energy detection does not suffice. Bytelight [45] proposes to provide indoor navigation to user's through visible light communication (VLC) from customized light bulbs. However, such customizations may not be possible in all applications. We refer readers to [46], for a rich list of different camera based localization techniques, for more details. In general, these approaches are limited in terms of range or require relatively large, costly and energy expensive cameras.

Augmented Reality Tagging. The most basic forms of AR tagging is in the form of QR codes and customized barcodes [47], [48]. Raskar et al. [49] propose an augmented reality system where users are notified of objects fit with passive RFID tags equipped with photoreceptors detecting light from the user's projection device. Similar to [50], Aoki et al. [51] propose an augmented reality system that send

context information, such as, location and ID, through IR devices to a camera feed – for processing or communication. The RFID readers, mobile projectors, and cameras are costly as well as energy intensive.

We note that our proposed synchronization approach to hybrid radio-optical beaconing significantly differs from existing multi-radio optimization techniques such as low-power wake-up [52] or intelligent switching between radios with different energy profiles [53], [54]. Indeed, existing wake-up techniques are complementary in that the RF link can also be used to wake up the IR transmitter when a receiver is detected in the IR range.

8 Conclusions

In this paper, we argue that a hybrid radio-optical beaconing approach can facilitate accurate and low-power recognition of objects within a user's view. Our approach leverages the high directionality characteristic of an infrared link for precise angle and distance estimation, and the low power nature of a radio link for synchronization and communication. The novelty of this design lies in the usage of a radio link to synchronize the infrared beacons such that very short high-energy infrared pulses could be used, which results in much reduced energy consumption, a simplified receiver design, and small hardware size. We prototyped the system by designing radio-optical tags and a wearable receiver, in the form of an object tracking eye-glasses. Our prototype receiver locates the infrared tags with an angular accuracy of 1.2 degree on the horizontal and vertical dimensions, and up to 9 m distance at very low battery power consumption, supporting tag battery life of the order of years. More importantly, the receiver is able to successfully recognize objects more than 97 percent of the time. We believe that with such accurate in-view recognition and long lifetime, our system can expand support from smart-glasses to also a wide range of applications in the wearable computing spectrum.

The main goal of the problem addressed in this paper was to achieve a *low-power* framework for *precise recognition* of contexts in a user's view-space. We designed a system where objects or contexts are tagged with transmitters that are used as active landmarks for recognition on a wearable device. We clearly understand that it not feasible to tag every object/context in the world and the deployment density of the tags will depend on the application. We feel this paper, through a novel design, contributes a step forward in low-power wearables' technology by designing low-power active landmarks or tags that assist recognition using wearable devices and that are operable for years.

Acknowledgments

This project has been funded by the US NSF under grant CNS 1065463. Ashwin Ashok and Chenren Xu are corresponding authors.

References

[1] (2013). Google Glass [Online]. Available: http://www.google.com/glass/start/
[2] (2008). Recon jet [Online]. Available: http://www.reconinstruments.com/products/jet/
[3] Meta [Online]. Available: https://www.spaceglasses.com/
[4] Microsoft Hololens [Online]. Available: https://www.microsoft.com/microsoft-hololens/en-us
[5] A. D. Cheok, X. Yang, Z. Z. Ying, M. Billinghurst, and H. Kato, "Touch-space: Mixed reality game space based on ubiquitous, tangible, and social computing," *Personal Ubiquitous Comput.*, vol. 6, nos. 5–6, pp. 430–442, 2002.
[6] Occulus vr – virtual reality headset for 3D gaming [Online]. Available: http://www.oculus.com/
[7] S. Rallapalli, A. Ganesan, K. Chintalapudi, V. Padmanabhan, and L. Qiu, "Enabling physical analytics in retail stores using smart glasses," in *Proc. 20th Annu. Int. Conf. Mobile Comput. Netw.*, 2014, pp. 115–126.
[8] (2013). Wikitude: The world's leading augmented reality SDK [Online]. Available: http://www.wikitude.com/
[9] H. Wang, S. Sen, A. Elgohary, M. Farid, M. Youssef, and R. R. Choudhury, "No need to war-drive: Unsupervised indoor localization," in *Proc. 10th Int. Conf. Mobile Syst., Appl. Serv.*, 2012, pp. 197–210.
[10] (2013). Inside Microsoft Research: Making purchases with zero effort [Online]. Available: http://tinyurl.com/bkjt9js
[11] L. Yang, Y. Chen, X.-Y. Li, C. Xiao, M. Li, and Y. Liu, "Tagoram: Real-time tracking of mobile RFID tags to high precision using cots devices," in *Proc. 20th Annu. Int. Conf. Mobile Comput. Network.*, 2014, pp. 237–248.
[12] F. Adib, Z. Kabelac, D. Katabi, and R. C. Miller, "3D tracking via body radio reflections," in *Proc. 11th USENIX Conf. Netw. Syst. Des. Implementation*, 2014, pp. 317–329.
[13] L. Zhang, X.-Y. Li, W. Huang, K. Liu, S. Zong, X. Jian, P. Feng, T. Jung, and Y. Liu, "It starts with igaze: Visual attention driven networking with smart glasses, "in *Proc. 20th Annu. Int. Conf. Mobile Comput. Netw.*, 2014, pp. 91–102.
[14] A. Mayberry, P. Hu, B. Marlin, C. Salthouse, and D. Ganesan, "Ishadow: Design of a wearable, real-time mobile gaze tracker," in *Proc. 12th Annu. Int. Conf. Mobile Syst. Appl. Serv.*, 2014, pp. 82–94.
[15] Infrared data association [Online]. Available: http://irdajp.info/
[16] Giga-ir high speed opto-communication [Online]. Available: http://www.google.com/glass/start/
[17] L. Korba, S. Elgazzar, and T. Welch, "Active infrared sensors for mobile robots," *IEEE Trans. Instrum. Measure.*, vol. 43, no. 2, pp. 283–287, Apr. 1994.
[18] G. Benet, F. Blanes, J. E Simó, and P. Pérez, "Map building using infrared sensors in mobile robots," in *New Developments in Robotics Research?* Commack, NY, USA: Nova, 2006.
[19] B. Firner, C. Xu, R. Howard, and Y. Zhang, "Multiple receiver strategies for minimizing packet loss in dense sensor networks," in *Proc. 11th ACM Int. Symp. Mobile Ad Hoc Netw. Comput.*, 2010, pp. 211–220.
[20] C. Yuodelis and A. Hendrickson, "A qualitative and quantitative analysis of the human fovea during development," *Vis. Res.*, vol. 26, no. 6, pp. 847–855, 1986.
[21] A. Ashok, C. Xu, T. Vu, M. Gruteser, R. Howard, Y. Zhang, N. Mandayam, W. Yuan, and K. Dana, "Bifocus: Using radio-optical beacons for an augmented reality search application," in *Proc. 11th Annu. Int. Conf. Mobile Syst. Appl. Serv.*, 2013, pp. 507–508.
[22] TSAL 5300 [Online]. Available: http://www.led-eshop.de/PDF/5mm/VISHAY_TSAL5300.pdf
[23] BPV10NF [Online]. Available: http://www.vishay.com/docs/81503/bp10nf.pdf
[24] R. I. Hartley and A. Zisserman, *Multiple View Geometry in Computer Vision*, 2nd ed. Cambridge, U.K.: Cambridge Univ. Press, 2004.
[25] New object recognition algorithm learns on the fly [Online]. Available: http://www.gizmag.com/learning-object-recognition-algorithm-byu/30512/
[26] G. Oberholzer, P. Sommer, and R. Wattenhofer, "Spiderbat: Augmenting wireless sensor networks with distance and angle information," in *Proc. 13th Annu. Int. Conf. Mobile Syst. Appl. Serv.*, 2011, pp. 241–256.
[27] R. LiKamWa, B. Priyantha, M. Philipose, L. Zhong, and P. Bahl, "Energy characterization and optimization of image sensing toward continuous mobile vision," in *Proc. 11th Annu. Int. Conf. Mobile Syst. Appl. Serv.*, 2013, pp. 69–82.
[28] J. Williams, 1991. Linear Technologies: High Speed Amplifier Techniques, A Designers Companion for Wideband Circuitry, http://cds.linear.com/docs/en/application-note/an47fa.pdf

[29] P. Bahl and V. N. Padmanabhan, "RADAR: An in-building RF-based user location and tracking system," in *IEEE Infocom*, vol. 2, pp. 775–784, 2000.
[30] K. Whitehouse, C. Karlof, and D. Culler, "A practical evaluation of radio signal strength for ranging-based localization," *ACM SIGMOBILE Mobile Comput. Commun. Rev.*, vol. 11, no. 1, pp. 41–52, 2007.
[31] M. Wattenhofer, "Lost in space or positioning in sensor networks," in *REALWSN*, 2005.
[32] B. Kusy, A. Ledeczi, and X. Koutsoukos, "Tracking mobile nodes using rf doppler shifts," in *Proc. 5th Int. Conf. Embedded Netw. Sensor Syst.*, 2007, pp. 29–42.
[33] J. Wang, F. Adib, R. Knepper, D. Katabi, and D. Rus, "RF-compass: Robot object manipulation using RFIDs," in *Proc. 19th Annu. Int. Conf. Mobile Comput. Netw.*, 2013, pp. 3–14.
[34] A. Harter, A. Hopper, P. Steggles, A. Ward, and P. Webster, "The anatomy of a context-aware application," in *Proc. 5th Annu. ACM/IEEE Int. Conf. Mobile Comput. Netw.*, 1999, pp. 59–68.
[35] N. B. Priyantha, A. Chakraborty, and H. Balakrishnan, "The cricket location-support system," in *Proc. 6th Annu. Int. Conf. Mobile Comput. Netw.*, 2000, pp. 32–43.
[36] A. Woo C. Karlof K. Whitehouse, F. Jiang, and D. Culler, "Sensor field localization: A deployment and empirical analysis," in UC Berkeley, Berkeley, CA, Tech. Rep. UCB//CSD-04-1349, 2004.
[37] A. Savvides, C.-C. Han, and Mani B. Strivastava, "Dynamic fine-grained localization in ad-hoc networks of sensors," in *Proc. 7th Annu. Int. Conf. Mobile Comput. Netw.*, 2001, pp. 166–179.
[38] E. Brassart, C. Pegard, and M. Mouaddib, "Localization using infrared beacons," *Robotica*, vol. 18, no. 2, pp. 153–161, 2000.
[39] M. V. Moreno, M. A. Zamora, J. Santa, and A. F. Skarmeta, "An indoor localization mechanism based on RFID and IR data in ambient intelligent environments," in *Proc. 6th Int. Conf. Innovative Mobile Internet Serv. Ubiquitous Comput.*, 2012, pp. 805–810.
[40] S. Lee and S.-Y. Jung, "Location awareness using angle-of-arrival based circular-pd-array for visible light communication," in *Proc. 18th Asia-Pac. Conf. Commun.*, 2012, pp. 480–485.
[41] R. Want, A. Hopper, V. Falcão, and J. Gibbons, "The active badge location system," *ACM Trans. Inform. Syst.*, vol. 10, no. 1, pp. 91–102, 1992.
[42] (1999). Firefly motion tracking system user's guide [Online]. Available: http://www.gesturecentral.com/firefly/FireflyUserGuide.pdf
[43] S. Hijikata, K. Terabayashi, and K. Umeda, "A simple indoor self-localization system using infrared leds," in *Proc. 6th Int. Conf. Netw. Sens. Syst.*, 2009, pp. 1–7.
[44] News - casio unveils prototype of visible light communication system [Online]. Available: bit.ly/zpfjY1
[45] Bytelight: Indoor location. with light [Online]. Available: http://www.bytelight.com/
[46] R. Mautz and S. Tilch, "Survey of optical indoor positioning systems," in *Proc. Int. Conf. Indoor Position. Indoor Navigation*, 2011, pp. 1–7.
[47] (2009). Artag [Online]. Available: http://www.artag.net/
[48] A. Mohan, G. Woo, S. Hiura, Q. Smithwick, and R. Raskar, "Bokode: Imperceptible visual tags for camera based interaction from a distance," *ACM Trans. Graph.*, New York, NY, USA, vol. 28, no. 3, pp. 98:1–98:8, Jul. 2009, Art. no. 98, Doi: 10.1145/1531326.1531404.
[49] R. Raskar, P. Beardsley, J. V. Baar, Y. Wang, P. Dietz, J. Lee, D. Leigh, and T. Willwacher, "RFIG lamps: Interacting with a self-describing world via photosensing wireless tags and projectors," in *ACM SIGGRAPH*, 2004, pp. 406–415.
[50] D. J. Moore, R. Want, B. L. Harrison, A. Gujar, and K. Fishkin, "Implementing phicons: Combining computer vision with infrared technology for interactive physical icons," in *Proc. 12th Annu. ACM Symp. User Interface Softw. Technol.*, 1999, pp. 67–68.
[51] H. Aoki and S. Matsushita, "Balloon tag: (In)visible marker which tells who's who," in *Proc. 4th IEEE Int. Symp. Wearable Comput.*, 2000, p. 181.
[52] E. Shih, P. Bahl, and M. J. Sinclair, "Wake on wireless: An event driven energy saving strategy for battery operated devices," in *Proc. 8th Annu. Int. Conf. Mobile Comput. Netw.*, 2002, pp. 160–171.
[53] G. Ananthanarayanan and I. Stoica, "Blue-fi: Enhancing wi-fi performance using bluetooth signals," in *Proc. 7th Int. Conf. Mobile Syst. Appl. Serv.*, 2009, pp. 249–262.
[54] M. Uddin and T. Nadeem, "A2PSM: Audio assisted wi-fi power saving mechanism for smart devices," in *Proc. 14th Workshop Mobile Comput. Syst. Appl.*, 2013.

Ashwin Ashok received the PhD degree in Oct 2014 from the Wireless Information Network Lab (WINLAB) at Rutgers University under the supervision of Profs. Marco Gruteser, Narayan Mandayam, and Kristin Dana. He is a postdoctoral research associate in ECE at Carnegie Mellon University (CMU) under the mentorship of Prof. Peter Steenkiste (CMU). He works in collaboration with Dr. Fan Bai (GM labs) and conducts research on building cloud computing systems for vehicles. His research work spans areas in mobile systems and computing, cloud computing, visible light and camera communications, vehicular networking, computer vision, augmented reality, wearable systems and localization. His thesis pioneered camera-based communication in mobile computing and visible light communication (VLC) through a novel concept called visual MIMO. During his doctoral tenure at Rutgers, he won the best teaching assistant award and also a research excellence award. He is currently the workshop co-chair for MobiCom 2016 and is also co-chairing the workshop on Wearable Systems and Applications (WearSys) at MobiSys 2016. He is an experimentalist and has keen interests in mobile system design and prototyping, and loves tinkering with circuits, Arduino and wearable devices.

Chenren Xu received the BE degree (Highest Hons.) from the Department of Automation at Shanghai University in 2008, the MS degree from the Department of Statistics at Rutgers in 2014, and the PhD degree under the guidance of Prof. Yanyong Zhang at WINLAB/Department of Electrical and Computer Engineering of Rutgers University in 2014. He is an assistant professor in the School of EECS at Peking University and a member of CECA. His research interests span different aspects of wireless networks. He is currently researching cross-layer design in the XIA future internet architecture to support and optimize mobility, with a particular focus on vehicular use cases. He is also interested in building health-related edge applications involving mobile sensing, ubiquitous computing, embedded systems and machine learning, to deploy them in the wild and benefit people in daily life.

Tam Vu received the BS degree in 2006 from Hanoi University of Technology, Vietnam, and the PhD degree in 2013 from WINLAB and the Department of Computer Science, Rutgers University, advised by Prof. Marco Gruteser and Prof. Dipankar Raychaudhuri. He is an assistant professor in the Department of Computer Science and Engineering, University of Colorado, Denver. He directs the Mobile and Networked Systems (MNS) lab, where he and his team have been working on security and privacy protections for mobile systems and context services. Collaborating with clinical doctors and and industrial partners, he is also leading research efforts to design and develop mobile healthcare systems that help monitor and diagnose sleeping disorders in children. Before that, he worked on various topics in core networking including inter domain routing, mobile-centric network architecture for future Internet, and end-host location privacy protection. He has interdisciplinary interests that combine algorithmic and theoretic aspects of computer science with hardware and software engineering, distributed systems and wireless communications of computer engineering. He is the recipient of CRC Interdisciplinary Fellowship in 2015. He received the Google Faculty Research Award in 2014, and best paper awards in ACM MobiCom 2012 and ACM MobiCom 2011. His research also received press coverage including *The Wall Street Journal*, National Public Radio (NPR), *MIT Technology Review*, Yahoo News, among other venues.

Marco Gruteser is a professor of electrical and computer engineering as well as computer science (by courtesy) at Rutgers University's Wireless Information Network Laboratory (WINLAB). He directs research in mobile computing, is a pioneer in the area of location privacy and recognized for his work on connected vehicles.

Rich Howard is a research professor at WINLAB, Rutgers University. He was formerly with Bell-Labs research as a director of Wireless Research.

Yanyong Zhang received the BS degree in computer science from the University of Science and Technology of China in 1997 and the PhD degree in computer science and engineering from Penn State University in 2002. She joined the Electrical and Computer Engineering Department of Rutgers University as an assistant professor in 2002. In 2008, she was promoted to associate professor with tenure, and in 2015, was promoted to full professor. She is also a member of the Wireless Information Networking Laboratory (Winlab). During March-July 2009, she was a visiting scientist at Nokia Research Center Beijing. Dr. Zhang is the recipient of the US NSF CAREER award. She is currently an associate editor for the *IEEE Transactions on Mobile Computing* and *IEEE Transactions on Services Computing*. She has served on TPC for many conferences, including Infocom, ICDCS, DSN, IPSN, etc.

Narayan Mandayam received the BTech (Hons.) degree in 1989 from the Indian Institute of Technology, Kharagpur, and the MS and PhD degrees in 1991 and 1994, respectively, from Rice University, Houston, TX, all in electrical engineering. He is a Distinguished Professor at Rutgers University. From 1994 to 1996, he was a research associate at the Wireless Information Network Laboratory (WINLAB), Dept. of Electrical & Computer Engineering, Rutgers University. In September 1996, he joined the faculty of the ECE department at Rutgers where he became associate professor in 2001 and professor in 2003. Currently, he also serves as associate director at WINLAB. He was a visiting faculty fellow in the Department of Electrical Engineering, Princeton University, in Fall 2002 and a visiting faculty at the Indian Institute of Science in Spring 2003.

Wenjia Yuan received the PhD degree from Rutgers University in Oct 2014. She is a software engineer with Google in Mountain View, California, USA.

Kristin Dana received the BS degree in 1990 from the Cooper Union (NY, NY), the MS degree from Massachusetts Institute of Technology in 1992, and the PhD degree from Columbia University (NY, NY) in 1999. She is an associate professor in the Department of Electrical and Computer Engineering at Rutgers, The State University of New Jersey. Her research is in computer vision, pattern recognition, machine learning, optics, and computer graphics. She is the inventor of the "texture camera" for convenient measurement of reflectance and texture. Dr. Dana is also a member of Rutgers Center for Cognitive Science, a member of Graduate Faculty of the Computer Science Department, and an adjunct assistant professor of medicine at UMDNJ, Robert Wood Johnson Medical School. From 1992-1995, she was on the research staff at Sarnoff Corporation developing real-time motion estimation algorithms for applications in defense, biomedicine, and entertainment industries. She is the recipient of the General Electric "Faculty of the Future" fellowship in 1990, the Sarnoff Corporation Technical Achievement Award in 1994 for the development of a practical algorithm for the real-time alignment of visible and infrared video images, and the 2001 US National Science Foundation Career Award for a program investigating surface science for vision and graphics.

▷ **For more information on this or any other computing topic, please visit our Digital Library at** www.computer.org/publications/dlib.