Hybrid Visible Light Communication for Cameras and Low-Power Embedded Devices

Niranjini Rajagopal Patrick Lazik Anthony Rowe
Electrical and Computer Engineering Department
Carnegie Mellon University
{niranjir,plazik,agr}@ece.cmu.edu

ABSTRACT

Visible light communication (VLC) between LED light bulbs and smart-phone cameras has already begun to gain traction for identification and indoor localization applications. To support detection by cameras, the frequencies and data rates are typically limited to below 1kHz and tens of bytes per second (Bps). In this paper, we present a technique for transmitting data from solid-state luminaries, used for interior ambient lighting, simultaneously to both cameras and low-power embedded devices in a manner that is imperceptible to occupants. This allows the camera communication VLC channel to also act as a higher speed downstream link and low-power wakeup mechanism for energy-constrained devices. Our approach uses Manchester encoding and Binary Frequency Shift Keying (BFSK) to modulate the high-speed data stream and applies duty-cycle adjustment to generate the slower camera communication signal. We explore the trade-off between the performance of the two communication channels. Our hybrid communication protocol is also compatible with existing IR receivers. This allows lights to communicate with low-cost commodity chipsets and control home appliances such as TVs, AV receivers, AC window units, etc. We show that we are able to reliably simultaneously transmit low-speed data at 1.3 Bps to camera enabled devices and higher-speed data at 104 Bps to low-power embedded devices. Since the majority of energy in many RF communication protocols often goes towards media access and receiving, VLC-triggered wakeup can significantly decrease system energy consumption. We also demonstrate a proof-of-concept wakeup circuit that consumes less then $204uA$ and can be triggered in less then 10ms.

Categories and Subject Descriptors

[**Networks**]: Wireless local area networks; [**Networks**]: Location based services; [**Information systems**]:

General Terms

Experimentation, Measurement

Keywords

Visible light communication, Hybrid data-rate, Camera communication

1. INTRODUCTION

Inside buildings, light bulbs are pervasive, have ample access to power, and are often well positioned for sensing and communication applications. VLC holds the promise of providing an extremely high data-rate and low-cost network link. VLC systems have the potential for high signal-to-noise ratios (greater than 50dB), are not regulated by the FCC, and can be contained easily within walls, providing a high degree of spatial diversity. In most high-rate VLC systems, the receiver is a device with a photodiode. Another common receiver for VLC is the camera in a mobile phone. This is particularly useful in localization of mobile phones where the lights act as landmarks or anchor points [1]. If each light can uniquely identify itself to nearby mobile devices, it would be possible to easily distinguish between rooms and areas within a room that are illuminated by different lights. Combining communication with geometric information such as the position of multiple lights in an image and angle-of-arrival, camera communication-based localization systems have achieved sub-meter accuracy [2]. Broadly, in the VLC space, we have high data-rate systems with photodiode-based receivers and lower data-rate systems with camera-based receivers.

In this paper, we investigate the problem of mixing camera communication with faster modulation designed for systems with photodiode receivers. This would allow LED bulbs to communicate with off-the-shelf mobile devices while simultaneously coordinating and acting as a downlink for ultra-low power tags and sensors.

Utilizing interior lights in this context for mixed communication is challenging for four main reasons: First, the system should be dimmable and flicker-free in order to provide high-quality indoor illumination. Second, since most smartphones no longer have high-speed optical sensors (like IR receivers), we would like to decode data using cameras that typically only operate at tens of hertz, while interleaving the higher-speed communication for low-power devices. Third, duty-cycled asynchronous low-power wakeup needs to coexist with the normal data streams. Finally, due to the dense

VLCS'14, September 7, 2014, Maui, Hawaii, USA.

http://dx.doi.org/10.1145/2643164.2643173.

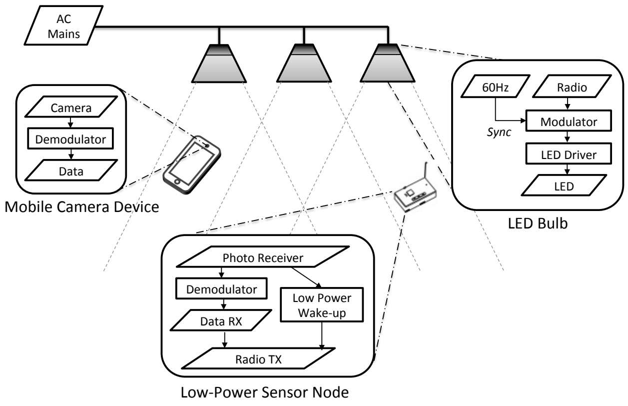

Figure 1: System architecture

configurations of indoor lighting, we require the ability to support scalable multiple-access among transmitters.

The main contribution of this paper is the design and evaluation of a protocol that communicates with low-power embedded devices and cameras simultaneously in a scalable manner for indoor applications.

2. SYSTEM ARCHITECTURE

Figure 1 shows our envisioned system architecture. We see LED lights that have a radio receiver and an onboard processor that is responsible for generating a modulated signal. The lights also have access to a common global 60Hz signal from the AC mains that can be used for time synchronization to help facilitate MAC coordination. Time synchronization can also be achieved using RF. Previous work shows how mobile devices like smart-phones can decode data at a low rate using their onboard camera (briefly reviewed in the next section). Low-power sensor nodes can leverage a high-speed signal for ultra-low power wakeup as well as downstream communication. All of the required hardware for the transmitters already exists in commercially available LED retrofits like the Philips Hue, LIFX, TCP connect, etc. We believe our solution should be compatible through a firmware modification to many of these devices.

2.1 Communicating with Low-Power Tags

When communicating with low-power embedded devices, we are primarily concerned with two modes of operation: First, there should be a readily accessible data channel for devices that wish to periodically wake up and check for control commands such as clock synchronization. Second, there should be an ability to asynchronously wake up devices in order to perform low-latency queries such as requesting for the last sampled sensor value. For example, in ultra-low power environments like those found in energy harvesting systems, one of the main practical concerns is that devices may only transmit data at a period of tens of minutes or hours. In many applications (like querying for a location), we would like to be able to immediately ping the device expending some of its energy budget when it is most valuable. This type of interaction can be enabled by an always-on VLC wakeup mechanism. We designed a proof-of-concept circuit that consumes less then $204uW$, but in a future revision this could be drastically decreased to below the drainage threshold of batteries. Our current implementation also does not operate over a wide dynamic range.

2.2 IR Receiver

There has been extensive work on photodiode-based VLC communication over the past few decades. Rather than inventing an entirely new low-level modulation scheme, we adapt existing work from IR communication. There are multiple common techniques used to transmit IR data that range in speed between 9.6kbps (SIR), all the way to 16Mbps (VFIR). We adopt an approach similar to the RC-5 standard that uses bursts of a particular frequency (37.9kHz) to modulate Manchester encoded data. In order to control the intensity of the signal for the camera communication, we can adjust the duty-cycle of the carrier bursts and pad non-carrier times with a similarly duty-cycled out-of-band frequency. Since overhead LED lights are extremely powerful compared to most IR LEDs in hand-held controllers, we are able to send IR data signals to existing appliances. Many of the new LED bulbs already have RF capabilities, enabling them to act as reliable proxies to replace older line-of-sight IR remotes.

3. RELATED WORK

There is a significant body of work for sending high-speed data over light using specialized hardware. The favorable characteristics of using unlicensed spectrum at bandwidths of up to 100MHz [3] with trichromatic (RGB) LEDs, or up to 20MHz with the more ubiquitous phosphorescent white LEDs [4, 5], make it an attractive contender for wireless communication. In 2006, the IEEE developed a draft VLC standard known as 802.15.7 [6]. All of these approaches use a variant of On-Off keying (OOK) for modulation. More complex schemes such as Quadrature Amplitude Modulation (QAM) and Discrete Multi-tone Modulation (DMT) / Orthogonal Frequency Division Multiplexing (OFDM) [7] are also possible, but require analog control over the LED intensity which increases hardware costs.

There have also been several efforts to use the rolling-shutter based smart-phone cameras as communication receivers. Danakis et. al. [8] exploit the rolling shutter effect of a smart-phone's CMOS camera to capture OOK modulated data from LEDs. The authors generate data sequences using Manchester encoding, resulting in multiple symbols being captured per frame. VRCodes (NewsFlash) [9] takes advantage of the rolling shutter effect to decode visual tags displayed on LCDs. The technology exploits the "flicker-fusion threshold" of the human eye to blend the tags into the background by rapidly flashing complimentary hues of color, still visible to a rolling shutter camera. CamCom [10] uses undersampled frequency shift OOK (UFSOOK) by encoding data at frequencies that are harmonics of the frame rate, and decoding data by processing the subsampled aliased frequencies. ByteLight [19] is a commercial effort for communication between overhead LED lights and mobile phones, by possibly exploiting the rolling shutter effect of the camera. In this paper, we will be adopting a rolling-shutter based BFSK camera communication method [1].

We believe that most the rolling-shutter based approaches would also be compatible with our hybrid approach of duty-cycle controlled high-frequency signals overlayed on top of lower rate communication schemes. In all cases, changing the duty-cycle will impact light and dark region contrast,

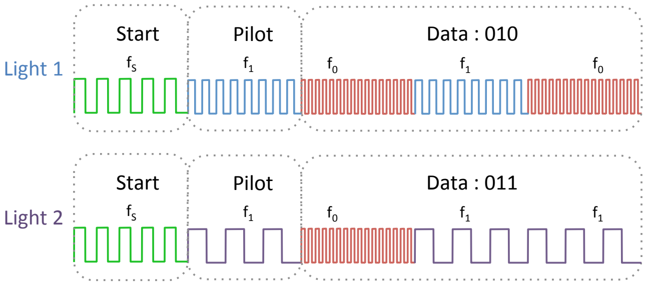

Figure 2: Camera communication modulation


(a) (b)
(c) (d)

Figure 3: Images as seen by a camera of a surface illuminated by an LED pulsing at (a) 1kHz (b) 100kHz. Mixed low-speed and high-speed communication with (c) 15-85% duty cycle (d) 40-60% duty-cycle

which will have some negative impact on performance. The main goal in this paper is to evaluate that impact and show that such an approach is feasible. To the best of our knowledge, we are one of the first to formally explore the topic of mixing camera communication with an infrastructure that communicates at higher speeds with low-power devices.

In [11] we see that it is possible to create a device-addressable low-power wakeup circuit in custom silicon that consumes as little as $695pW$ of power. While a custom silicon design probes many of the fundamental limits of low-power wakeup, we provide a simpler mechanism using off-the-shelf discrete components as a simple proof of concept for battery operated devices like wireless sensor nodes.

4. COMMUNICATION SCHEME

In this section, we present an overview of the camera communication scheme, followed by our proposed scheme of mixing the high-rate data stream with the camera communication data stream.

4.1 Camera Communication

Figure 2 shows an outline of our approach, which uses Binary Frequency Shift Keying (BFSK) to modulate the signal driving the light [1]. A preamble indicates the start of each data packet, followed by a pilot which is identical to a transmitter's *on* symbol. The pilot allows the receivers to measure the noise floor of each transmission. Multiple transmitters co-exist by using Frequency Division Multiple Access (FDMA), where each is allocated a unique frequency for transmitting its *on* symbol.

The receiver is a CMOS rolling shutter camera that captures the signal either through direct LOS or reflected from surfaces. Rolling shutter sensors expose and read out individual rows of pixels in a pipelined fashion. As a result, an LED pulsing at a frequency produces a banded image, the frequency of which is proportional to the frequency of the LED.

After a sequence of image frames are acquired, the spatial variation of the frequency components over the images is used to estimate the input frequency of the lights. The demodulation and decoding is performed in software to detect the received IDs. With a frame rate of 30 frames per second, a symbol duration of 33ms, and a frame length 14 bits, a byte of data takes $462ms$ to transmit. It is possible to decode the camera communication signal using a photodiode in the same manner. However, for a low-power tag, the energy consumed by staying on and processing data for this amount of time is not trivial. The limiting factor for the low data rate of the camera communication is the frame rate of the camera and the rolling shutter scan rate. The photodiode receiver need not be constrained by these timings. This motivates the need for a scheme which can simultaneously communicate with phones and tags at two different data rates (two different channels) - but share the same physical channel.

4.2 Incorporating High-Rate Data

When overlaying high-speed data, we assume an LED driver that generates a binary output, which is most common in low-cost LED bulbs. We use Manchester encoding with BFSK modulation, where a burst of pulses of two different frequencies represent the high and low logic levels of the encoded data.

Figure 3(a) and Figure 3(b) show simulated images of what a camera would capture in the presence of a light pulsing at 1kHz and 100kHz respectively. The column scan time is 47us, which is approximately the same as a camera that captures 1080p frames at 30fps (iPad 3). The image in Figure 3(a) tends to look gray because the column scan time is equivalent to 4.7 cycles of the 100kHz signal causing the average intensity of the light in this duration to be captured uniformly across the entire row. On the other hand, the bands are prominent in Figure 3(b) since every *on* or *off* period of the light at 1kHz spans across approximately 10 columns of the image.

The frequency of the LED is determined by the high-speed data and the duty-cycle is determined by the low-speed data. When the high-speed modulated data overlaps with the *on* time of the light for the low-speed modulated data, the pulses are duty-cycled such that the *on* time is longer than the *off* time. Similarly, when the high-speed modulated data overlaps with the *off* time of the light for the low-speed modulated data, the pulses are duty-cycled such that the *off* time is longer than the *on* time. This results in light gray and dark gray stripes instead of white and black stripes on the images.

Figure 3(c) and Figure 3(d) show images that we would expect to see by varying the duty cycle of the 100kHz signal based on overlapping the 1kHz signal of the low-speed communication. Figure 3(a) is equivalent to the mixed communication system with 50-50% duty cycle and Figure 3(b) is equivalent to the mixed communication system with 0-100% duty cycle. Figure 4 shows a snapshot of the proposed modulation scheme. The top row shows the BFSK modulated low speed data as proposed in our previous work [1]. The second row shows the Manchester encoded high speed data. The third row shows the modulated data where the frequency is determined by the high-speed encoded data

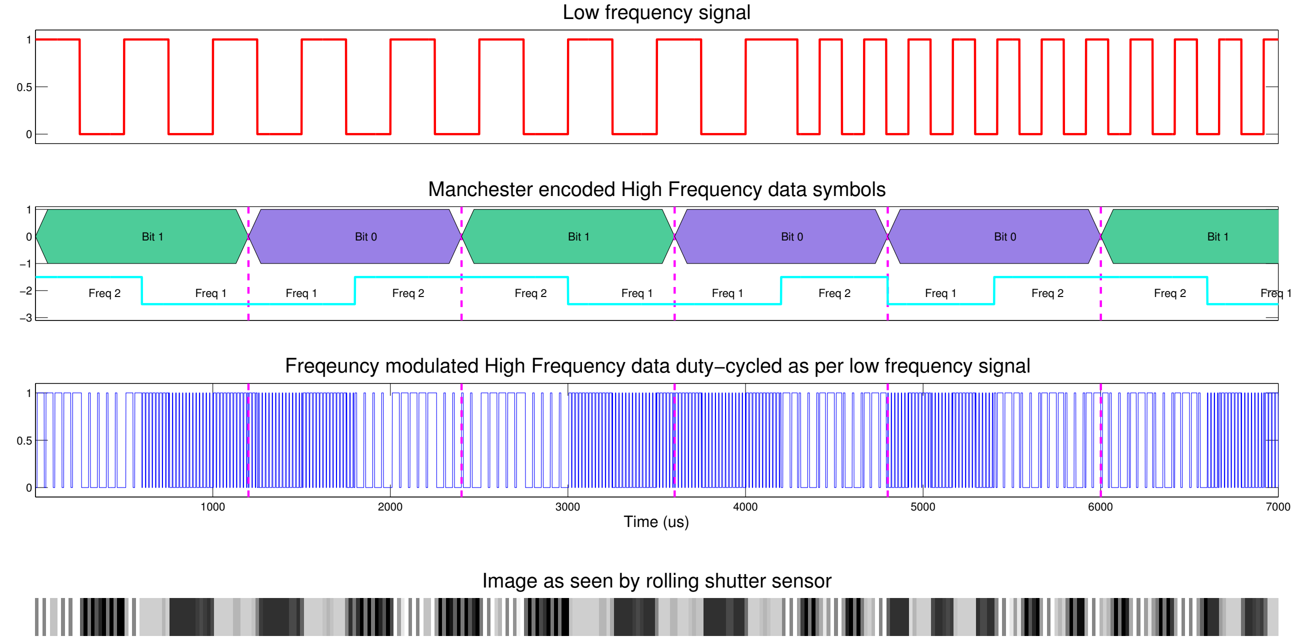

Figure 4: Mixed low-speed and high-speed communication

and the duty-cycle is determined by the low-speed data. The bottom row shows the resulting image from the camera.

The receiver for the low-speed data can be any smartphone with a rolling-shutter based camera. To receive the high-speed data, we can use a photodiode and demodulate the BFSK data either by analog frequency detection (often performed in hardware) or performing the frequency analysis after acquiring the data. The thresholded signal is then decoded by a Manchester decoder. We propose using an IR Receiver Module (for example a RPM7138-R) for the receiver. This is commonly used for remote controls and has a reasonably low current consumption of 0.95mA. This chip is designed for an IR transmitter at 37.9kHz. In our scheme, we use a burst of pulses of frequency 37.9Hz and an out-of-band frequency (10kHz) for modulating the Manchester encoded data to avoid flicker. This results in a high-speed data rate of 104Bps. The limiting factor with respect to data rate is the IR receiver, however, other receivers can operate at higher rates.

4.3 Multiple Access

The camera communication uses assigned frequencies to support multiple transmitters within a single collision domain through FDMA. Since the embedded tags are not able to simultaneously decode multiple frequencies, we suggest using a Time Division Multiple Access (TDMA) scheme to separate transmissions in time. Though not the focus of this work, TDMA time slots can be picked based on the FDMA frequencies used by each light.

5. LOW-POWER WAKEUP

Our low-power wakeup circuit is activated by receiving an unmodulated 100 kHz signal, broadcast from a VLC transmitter. A photodiode captures all incoming light, which is passed through multiple high-pass filters. The filters attenuate any out-of-band data signals and trigger a low-power comparator if a wakeup signal is received. The comparator digitizes the received signal, which may then be used to wake an attached microcontroller. Due to the limited bandwidth-gain-product of the op-amps in the filters, the circuit uses two cascaded op-amps to provide a relatively flat frequency-response and hence equal amplification to more than 100kHz.

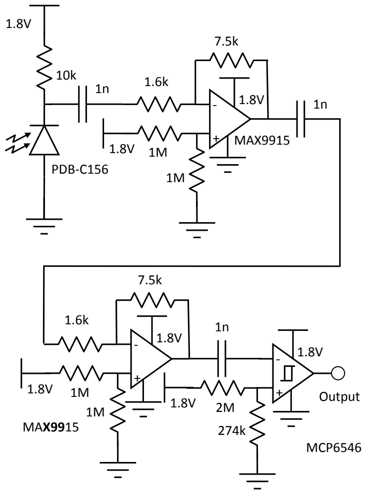

Figure 5: Proof-of-concept wakeup circuit

The circuit consumes approximately 204μA at 1.8V when active and may be further duty-cycled to reduce power. The threshold on the comparator is tunable to allow for different light intensities. Though functional, this circuit exhibits relatively poor dynamic range. We intend to refine and further evaluate how to build efficient wakeup circuits, but this design confirms feasibility.

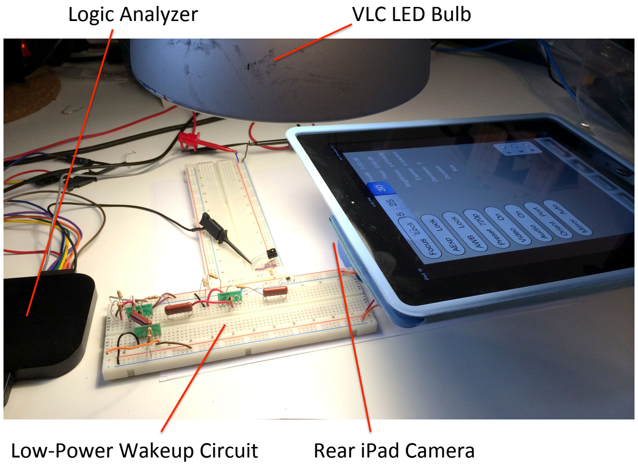

Figure 6: Experimental setup

6. EVALUATION

Our experimental setup is shown in Figure 6. We based our prototype light on a 9.5 Watt (60W incandescent equivalent) Cree warm white (2700K) LED bulb, which outputs 800 lumens and is currently available for less than $9. We replaced the bulb's power electronics with a simple MOSFET driver circuit, which is controlled by a function generator sending it synthesized signals. We use an adjustable aperture over the light for controlling the illumination. The receiver for the high-speed data is a standard infrared receiver module, commonly used in remote control receivers (ROHM RPM7138-R), connected to a logic analyzer. We use the front-facing camera on an iPad 3 collecting 720p HD video at 30fps as a receiver for the low-speed data. It is mounted 5cm above a white surface. We generate the signals in MATLAB and transmit it to the LED using a programmable signal generator. Captured images are then sent back to MATLAB for processing.

6.1 Base-Line Performance Variation With Illumination

With high speed data modulated at around 40kHz and support for low duty-cycle (10%), a sampling frequency of at least 400kHz is required. A symbol length of 33ms (one frame at 30fps) for the camera communication restricted the data to an equivalent of three symbols of low-speed data due to the limited buffer size in our function generator. Since this was insufficient to send a full length data sequence, we first obtained the BER-SNR mapping (Figure 7(a)) for the camera communication without the high-speed data by transmitting 25 random packets consisting of 6 bits of low speed data. For all further tests with high-speed data, we transmit 25 random packets consisting of 16 bits of high speed data and 3 symbols of low speed data. In these tests, we evaluate the BER of high speed data and SNR of the camera communication.

Figure 7(b) shows the performance of both channels at 20% duty cycle over a range of transmission luminosities. In our experiment the system was able to decode data correctly down to a transmission illumination of 220lux, after which the BER increases rapidly. The SNR of the low-speed signal increases exponentially over the measured illuminations.

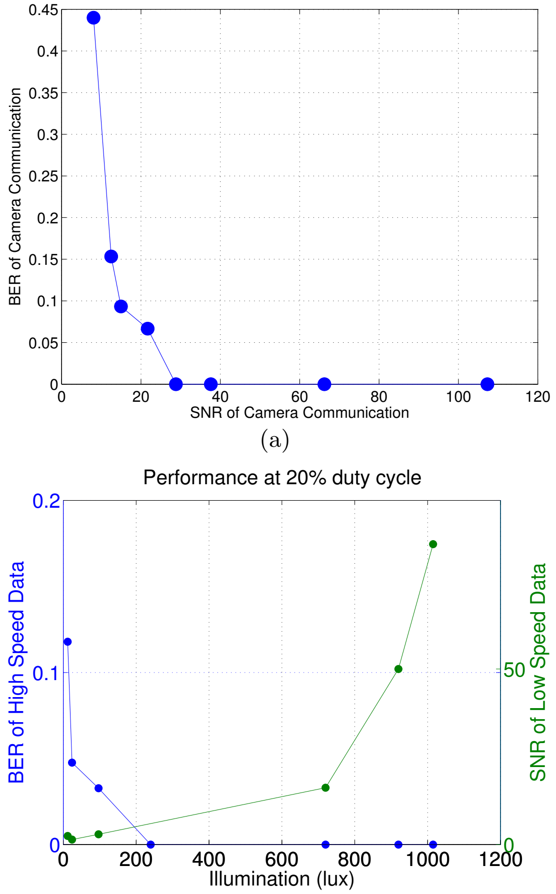

(a)

(b)

Figure 7: (a) BER to SNR mapping for camera (b) Performance of both channels with change in illumination

6.2 Trade-Off Between High and Low-Rate Communication

Figure 8(a) shows the impact of the high-speed data duty-cycle on both communication channels at 12lux of luminosity, measured by a light meter on the white surface in the field of view of the iPad's camera. We see an inverse relationship between the duty-cycle of the high-speed data and the BER, as well as an inverse relationship with the SNR of the low-speed component. This is to be expected since the high-speed component reduces the contrast of low-speed data. Figure 8(b) shows the same experiment at an illumination of 720lux. With more light, we see that the BER of the high-speed data is vastly reduced due to the increase in transmission power and the SNR of the low-speed component is approximately doubled for each duty-cycle setting. The illumination value of this setup corresponds to what can be expected in a normal office environment.

6.3 Future Evaluation

There are a few more trade-offs that we would like to evaluate in the future, which are currently not possible since our IR receiver fixes the modulation frequency to 37.9kHz (see Section 2.1).

There is a trade-off between the performance of both communication systems as the ratio of high-speed frequency to

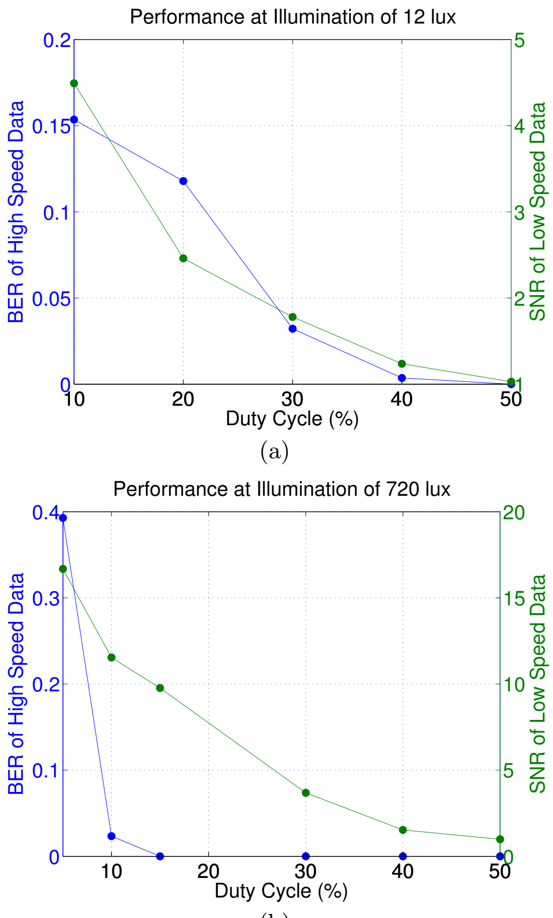

(a)

(b)

Figure 8: Trade-off in performance between high and low-speed channels with change in duty-cycle in (a) low-light (b) bright light

low-speed frequency increases. At a fixed low-speed frequency, as we increase the high-speed frequency, the performance of the low-speed channel would improve due to increased smoothing of the high-speed data by the rolling shutter. However, the high-speed channel's performance would degrade due to lesser samples per symbol.

There is also an interesting trade-off between the high-speed data rate and the SNR of the camera communication as we vary the duty-cycle. As we increase the duty-cycle of the modulated data (for example from 40%-60% to 20%-80%), the SNR of the camera communication would improve (see Section 6.2). However, the maximum possible data-rate of the high-speed data would reduce if we sample at the Nyquist frequency.

7. CONCLUSIONS AND FUTURE WORK

In conclusion, this paper provides a scheme for mixing camera communication with higher-speed data through use of adaptive duty-cycling. We see that there is an inherent trade-off between the performance of the two schemes. More extreme duty-cycles in the high-speed data improve the camera communication, but decrease the high-speed performance. We also show a wakeup circuit that can be used to support aggressive duty-cycling of low-power embedded devices like wireless sensors. In the future, we intend to investigate how this approach can be used to support more complex sensing applications as well as services such as time synchronization. We also believe we can significantly improve our proof-of-concept low-power wakeup circuit. We intend to redesign the receiver without the data-rate restriction imposed by the IR receiver, and evaluate the system at higher data-rates.

8. ACKNOWLEDGEMENTS

This research was funded in part by TerraSwarm, one of six centers of STARnet, a Semiconductor Research Corporation program sponsored by MARCO and DARPA.

9. REFERENCES

[1] N. Rajagopal, P. Lazik, and A Rowe. Visual light landmarks for mobile devices. In *Proceedings of the 13th ACM/IEEE Conference on Information Processing on Sensor Networks*, 2014.

[2] Y. Kuo, P. Pannuto, K. Hsiao, and P. Dutta. Luxapose: Indoor positioning with mobile phones and visible light. In *Proceedings of the 20th International Conference on Mobile Computing and Networking*, 2014.

[3] T. Komine and M. Nakagawa. Fundamental analysis for visible-light communication system using led lights. *IEEE Transactions on Consumer Electronics*, 50(1):100–107, 2004.

[4] Jelena Grubor, Sian Chong Jeffrey Lee, Klaus-Dieter Langer, Ton Koonen, and Joachim W. Walewski. Wireless high-speed data transmission with phosphorescent white-light leds. In *33rd European Conference and Exhibition of Optical Communication - Post-Deadline Papers*, pages 1–2, 2007.

[5] J. Grubor, O. C. Gaete Jamett, J. W. Walewski, and K. d. Langer. High-speed wireless indoor communication via visible light. In *ITG Fachbericht*, volume 198, pages 203–208, sept. 2007.

[6] Sg vlc project draft 5c. *IEEE Draft Std. IEEE P802. 15-08-0667-01-0vlc*, September 2008.

[7] H. Elgala, R. Mesleh, and H. Haas. Indoor broadcasting via white leds and ofdm. *IEEE Transactions on Consumer Electronics*, 55(3):1127–1134, 2009.

[8] Christos Danakis, Mostafa Afgani, Gordon Povey, Ian Underwood, and Harald Haas. Using a cmos camera sensor for visible light communication. In *IEEE Globecom Workshops*, pages 1244–1248, 2012.

[9] G. Woo, A. Lippman, and R. Raskar. Vrcodes: Unobtrusive and active visual codes for interaction by exploiting rolling shutter. In *IEEE International Symposium on Mixed and Augmented Reality*, ISMAR '12, pages 59–64, 2012.

[10] Richard D. Roberts. Undersampled frequency shift on-off keying (ufsook) for camera communications (camcom). WOCC '13, 2013.

[11] Gyouho Kim, Yoonmyung Lee, Suyoung Bang, Inhee Lee, Yejoong Kim, D Sylvester, and D. Blaauw. A 695 pw standby power optical wake-up receiver for wireless sensor nodes. In *Custom Integrated Circuits Conference (CICC), 2012 IEEE*, pages 1–4, Sept 2012.