CeilingTalk: Lightweight Indoor Broadcast Through LED-Camera Communication

Yanbing Yang*, Jie Hao*, and Jun Luo

Abstract—Although Visible Light Communication (VLC) is gaining increasing attentions in research, developing a practical VLC system to harness its immediate benefits using Commercial Off-The-Shelf (COTS) devices is still an open issue. To this end, we develop and deploy CeilingTalk as a lightweight wireless broadcast system using COTS LED luminaries as transmitters and smartphone cameras as receivers so that it can be fully hosted in a smartphone and is feasible for all possible indoor environments. CeilingTalk innovates in both encoding and decoding to achieve an adequate throughput for realistic applications. On one hand, it employs Raptor coding to allow multiple LED luminaries to transmit collaboratively so as to benefit both throughput and reliability. On the other hand, it involves a lightweight decoding scheme to handle the asynchrony (both spatial and temporal) in transmissions. Moreover, we analyze the impact of various parameters on the performance of CeilingTalk, in order to derive a model for such VLC systems enabled by COTS devices and hence provide general guidance for future VLC deployments in larger scales. Finally, we conduct extensive field experiments to validate the effectiveness of our LED-camera VLC model, as well as to demonstrate the promising performance of CeilingTalk: up to 1.0 kb/s at a distance of 5 m.

Index Terms—Visible Light Communication, Raptor Code, Mobile Computing.

1 Introduction

Using visible light spectrum as the communication media has a long history that can be traced back to 19th century, and Visible Light Communication (VLC) has been holding the promise of extremely high data-rate and very low-cost for a few decades [2]. Nevertheless, it is the boosting market of Light Emitting Diode (LED) that has drastically accelerated the development of VLC as a supplement to existing wireless standards (e.g., Wi-Fi). As LEDs are becoming the pervasive lighting source for indoor environments such as shopping malls and office/residence buildings, and they have a ready access to power/information networks, we now have a handy infrastructure for implementing VLC-enabled wireless systems. However, whereas major research efforts have been made to boost the capacity of VLC [3], [4] and to enable indoor localization [5]–[7], harnessing the potential of LED-VLC using Commercial Off-The-Shelf (COTS) devices seems to be largely neglected.

Following the seminal work of [8], implementing a receiver applying the rolling-shutter effect of a CMOS camera has become a de facto standard for LED-VLC using COTS devices [6], [9]: only a customized light sensor may serve as an alternative because light sensors in smartphones are optimized for dynamic range rather than response speed [10]. So it appears to be a common belief that enabling LED-VLC using COTS devices is an addressed problem. However, the system built in [8] requires a secondary medium, a plain surface (to avoid interference from textures), for light reflecting. While the Signal Noise Ratio (SNR) of such reflected VLC could be very low in a well-lit room, the difficulty in finding a satisfactory surface in an indoor space may substantially confine the practicality of such a system. Consequently, the validity of the derived models and insights also become questionable.

In fact, VLC has quite a few advantages over existing wireless infrastructures (e.g., Wi-Fi and BLE) even without its high data-rate promise. First, the communication infrastructure is virtually free due to the default lighting requirement of indoor spaces. Secondly, the energy consumption of LED transmission is negligibly low as it is incurred only by the control units. Last but not least, the location-bound communication ability of VLC meets perfectly the need from the long envisioned location-based service: coupons or advertisements can be delivered exactly according to the users' locations. Practical VLC implementations using COTS devices may immediately bring these benefits to us, while allowing for a better understanding of future VLC deployments from theoretical and modeling perspectives.

In this paper, we re-visit the issue of realizing VLC-based broadcast using LEDs as transmitter and smartphone cameras as receiver. On one hand, we aim to build a practical system that works perfectly under normal lighting conditions. We design a coding mechanism that enables cooperative transmission among multiple close-by light sources, and we also innovate in an efficient decoding scheme to handle the transmission asynchrony. On the other hand, we intend to extract model parameters (e.g., which factors affect the data rate) and hence provide guidelines on future deployments of VLC system. Our major contributions are:

---

* *Equal contributing authors.* *Preliminary results were presented in Proceedings of IEEE INFOCOM 2016 [1].*

• *Yanbing Yang and Jun Luo are with the School of Computer Science and Engineering, Nanyang Technological University (NTU), Singapore. Email: {yyang017, junluo}@ntu.edu.sg.* *This work was supported in part by the AcRF Tier 2 Grant MOE2016-T2-2-022.*

• *Jie Hao is with the College of Computer Science and Technology, Nanjing University of Aeronautics and Astronautics, China. Email: haojie@nuaa.edu.cn.* *This work was started when she worked as a post-doctoral researcher in NTU.*

- We build CeilingTalk as a practical wireless broadcast system, using common LEDs as transmitters and smartphone cameras as receivers, without relying on media other than the light itself.
- We are the first to employ rateless codes to enhance reliability and also to improve the network throughput through cooperative transmissions.
- We also innovate in a lightweight decoding scheme that avoids intensive image processing for tackling both spatial and temporal asynchrony in transmissions, making it amenable to a full smartphone-based implementation.
- We develop a communication model to characterize this LED-camera VLC by extracting parameters from the experiments on CeilingTalk.
- We conduct extensive field experiments to validate the effectiveness of our LED-camera VLC model, as well as to demonstrate the promising performance of CeilingTalk.

The rest of our paper is organized as follows. We first survey related literature and introduce the preliminary of LED-camera communication in Sec. 2. Then we illustrate the system architecture of CeilingTalk in Sec. 3 and analyze the factors impacting system performance in Sec. 4. The extensive evaluations on CeilingTalk are reported in Sec. 5. Finally Sec. 6 concludes this paper.

2 Related Work and Background

VLC has been favored as a complementary approach to traditional wireless communications over the last decade, as it has the potential to offer broadband communication on unlicensed spectrum with a high degree of space reuse. The recent standard 802.15.7 [2] has defined specifications while categorizing VLC into two classes: high rate and low rate. High rate VLC can achieve up to gigabits per second with specialized high speed photodiode receivers [3], [11]–[13], whereas low rate VLC is generally explored on COTS mobile devices [8], [14]. As we focus only on the latter in this paper, so are the following discussions on the related literature and the preliminary about rolling shutter effect.

2.1 Literature Review

LED-Camera Communication mostly exploits the rolling-shutter effect of CMOS camera on smartphones. Unlike the high rate VLC where similar LED transmitters are used, the bottleneck lies on the smartphone receiver, which is where innovations can be made. Recent research [8] achieves data rates up to 148 b/s at 20fps. This low rate can be attributed to its low SNR resulted from the need for a second media, i.e. a plain surface, to reflect light so as for the camera to capture the banded images. As an important application of this type of VLC, Visible Light Positioning (VLP) [6], [7], [9] makes use of light as the only communication media to transmit location identifiers, and this is further extended by RollingLight [14] to build a new LED-camera VLC system. Color shift keying (CSK) is employed in ColorBars [15] to improve the throughput by utilizing tri-RGB LED luminaires. Nevertheless, all these proposals require either complicated modulation or heavy decoding computations that may not be feasible for COTS devices to operate in a long distance, and they treat several transmissions independently rather than leveraging them for throughput/reliability enhancement. Our CeilingTalk is designed to fill these gaps so as to make VLC practical, and it significantly outperforms the state-of-the-art RollingLight [14] and ColorBars [15] in either throughput or transmission distance. We summarize these proposals in TABLE 1.

TABLE 1
Summary of LED-Camera VLC Systems.

	Capacity	Distance (m)	Error Correction
Seminal [8]	148 b/s	0.09	N.A.
Luxapose [7]	N.A.	6.00	N.A.
VLandmark [6]	1.25 B/s	6.50	Hamming codes
HybridVLC [9]	1.30 B/s	0.05	N.A.
RollingLight [14]	11.32 B/s	1.60	Parity
ColorBars [15]	5.2kb/s	0.03	Reed-Solomon codes
CeilingTalk	1.0 kb/s	5.00	Raptor codes

As another important branch of low rate VLC, *Screen-Camera Communication* mainly focuses on designing sophisticated coded images to boost data rate [16]–[19], or to enhance link reliability [20]–[22]. To improve data rate, SBVLCD [16] and COBRA [17] exploit advanced barcode design, PixNet [18] leverages an efficient modulation mechanism, and SoftLight [19] employs channel coding. Reliability is addressed in terms of either frame synchronization [22] or barcode detection under a prolonged communication distance and device diversity [20], [21]. Moreover, recent studies intend to enhance viewing experience by hiding information in a given screen content without interfering communications. PiCode [23] integrates barcodes with existing images to reduce the visual artifacts, while HiLight [24] avoids modifications on RGB values by utilizing the alpha channel to encode bits into the pixel translucency changes. In most cases, Screen-Camera VLC only offers a sub-meter transmission distance (unless a large screen is dedicated as the transmitter [18]), as briefly summarized in TABLE 2, thus substantially confining its application domain.

TABLE 2
Summary of Screen-Camera VLC Systems.

	Capacity	Distance (m)	Modulation/Coding
SBVLC [16]	200 kb/s	0.17	Barcode
COBRA [17]	172 kb/s	0.25	Barcode
PixNet [18]	12 Mb/s	10	OFDM
SoftLight [19]	317.3 kb/s	0.20	Rateless codes
HiLight [24]	1.1 kb/s	0.27	BFSK

2.2 Rolling-Shutter Effect

Since the rolling-shutter effect of a CMOS camera is an important building block of CeilingTalk, we briefly introduce its basics along with related terminologies before diving into our system design. Suppose we have an LED luminaire emitting modulated light (i.e., it turns ON and OFF according to a preset pattern) that is to be received by a CMOS camera. The rolling-shutter effect of a camera, instead of exposing the whole image at once, conducts the exposure

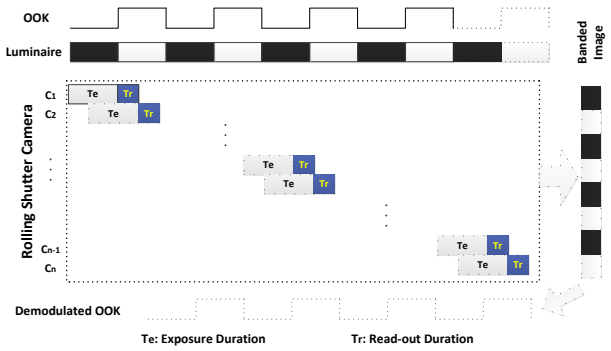

Fig. 1. Rolling-shutter effect of a CMOS camera on a smartphone.

in a column-by-column (or row-by-row) manner: a column is exposed for T_e before it is read out for T_r, and the next column gets exposed T_r after the previous one. When the luminaire is ON and a column is exposed at the same time, the column is bright; otherwise it is dark. As a result, the camera is sampling the intensity modulated signals with an interval T_r and thus a sampling frequency $F_r = 1/T_r$, leading to a banded image that carries the information sent by the luminaire. Here F_r is known as *rolling-shutter frequency*; it actually bounds the modulation frequency from above: as F_r is often in the order of tens of kHz, the supported modulation frequency is up to several kHz according to Nyquist Sampling Theorem [25]. A graphical illustration is provided in Fig. 1. In reality, the camera may expose and read out multiple columns at the same time, resulting in another important parameter W_r, or *rolling-shutter width*. Consequently, in a banded image sampled by the camera, either the whole W_r columns are bright or they are dark; this further suggests that the width of each band should be a multiple of W_r. The detailed relations between these parameters and the throughput shall be explored in 4.2.

3 System Architecture

CeilingTalk consists of two main components: COTS LED luminaires as stationary transmitters to emit intensity modulated light and smartphones as receivers. A user holding a smartphone may freely wander in the indoor environment and receive location-bound information from nearby LED luminaires, as illustrated in Fig. 2. In this section, we focus on presenting our innovations in designing the coding/decoding mechanisms. In particular, we explain

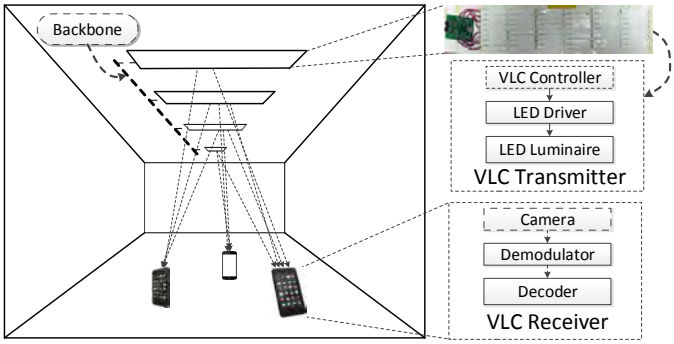

Fig. 2. System architecture of CeilingTalk.

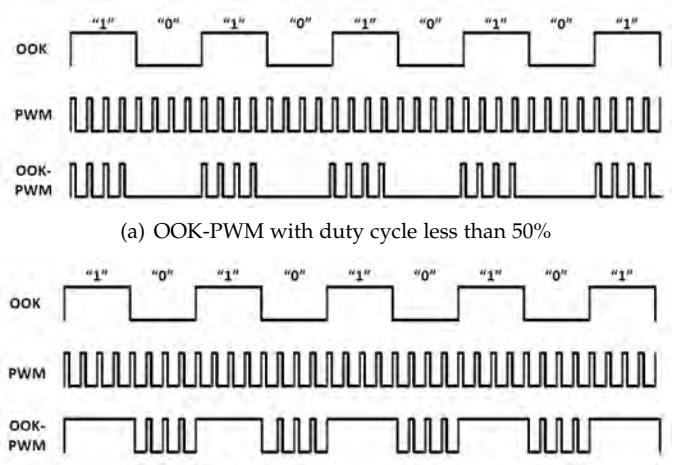

(a) OOK-PWM with duty cycle less than 50%

(b) OOK-PWM with duty cycle more than 50%

Fig. 3. OOK-PWM modulation: the combination of OOK-PWM effectively avoids flicker and also allows for dimming control.

how we may combine several transmissions from multiple luminaires to achieve higher throughput and reliability. And we also describe an efficient decoding scheme that avoids computation intensive image processing for handling both spatial and temporal asynchrony in transmissions.

3.1 A Hybrid OOK-PWM Modulation

CeilingTalk employs a hybrid modulation of On-Off Keying (OOK) and Pulse-Width Modulation (PWM) as shown in Fig. 3. As the simplest modulation method, OOK is universally employed on VLC, where the LED is switched ON or OFF to represent data bits "1" or "0", respectively. We further embed a high frequency PWM in each OOK bit for dimming control: it allows the light intensity of an LED luminaire to be adjusted without either causing flicker or interfering data transmission. We refer to Sec. 3.5 for more details. As the PWM is supposed to be undetectable by the rolling-shutter effect at the receiver, it runs at 80kHz in CeilingTalk.

The packet structure in CeilingTalk is set according to Fig. 4. Each packet has a preamble of bit string "111110" and it always ends with a single bit "0". The Packet Sequence Number (PSN) comes after the preamble and it is followed by the payload. For bits in both the PSN and payload, CeilingTalk applies 4B6B run length limit (RLL) line coding [2] for its simplicity and absent of DC component. As 4B6B coding guarantees no more than four adjacent "1" or "0" bits, the preamble generates the widest bright band in the image and is hence easy to be detected as the indication of the start of a packet. Note that the five adjacent "1" in the preamble would not cause flicker. Within the duration for one packet, there are only 5 continuous bright bands out of 67 bands in total. The flicker factor is thus too low to be

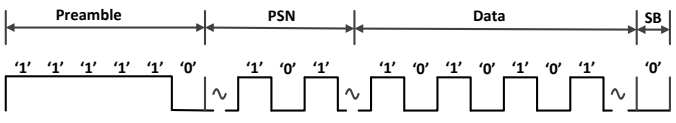

Fig. 4. CeilingTalk packet structure: a packet consists of a preamble "111110", up to 8-bit PSN, 32-bit data payload, and a tailer "0".

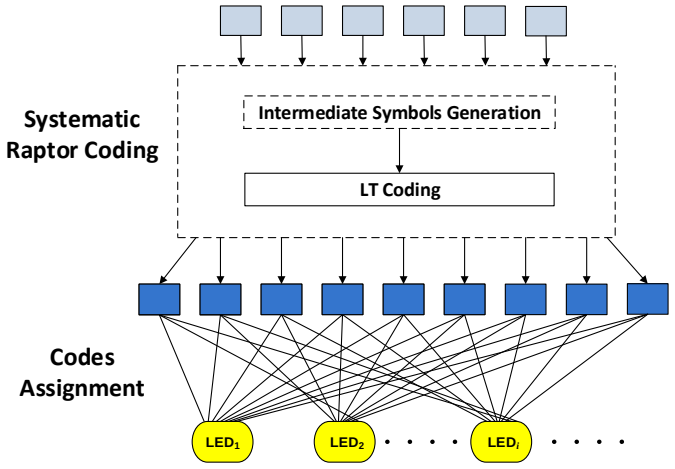

Fig. 5. Raptor coding and codes assignment.

discernable. Moreover, we can adjust the duty cycle of PWM for the continuous bands to further alleviate flicker.

3.2 Encoding and Code Assignment

As LED-Camera VLC is normally unidirectional, a transmitter has no knowledge about the reception status of a receiver. Therefore, a Forward Error Correction (FEC) scheme has to be in place to combat the packet loss. Moreover, simultaneous transmissions from multiple luminaires also call for an encoding mechanism that enables the transmitters to cooperatively improve the throughput and reliability. To this end, rateless codes appear to be an appealing choice, as it encodes k original packets into potentially infinite number of packets so that the receiver can successfully recover the original packets by receiving any $m > k$ encoded packets.

Among all rateless codes, Raptor code [26] comes with linear time encoding and decoding, causing low computation complexity and decoding overhead. Therefore, CeilingTalk adopts Raptor code and its implementation [27]. As shown in Fig. 5, a set of intermediate packets are firstly derived from the original packets so that the intermediate packets can sufficiently reconstruct the original ones. Repair packets are then produced by applying LT encoding; each is derived by XORing a number of intermediate packets. The final encoded packets are the combination of the original and repair ones. As CeilingTalk has no acknowledgement from camera for the transmitter to stop coding and sending, we regulate the transmitter to send out only $n = 1.25k$ encoded packets for every k original packets. Upon receiving slightly more k packets, a receiver uses Gaussian elimination to start recovering the original ones.

As CeilinTalk relies on the lighting system to transmit packets, it has the opportunity to leverage multiple close-by luminaires to perform cooperative transmissions for improving both throughput and reliability. Imagine a simple scenario that a smartphone receives messages from two luminaires. Suppose we let the luminaires to work independently, then either we waste the capacity of one luminaire if the message is sent by only one luminaire or individual packets loss may ruin both messages if two luminaires are transmitting. However, if we allow the two luminaires to cooperatively transmitting, then we may either reduce the transmission time or let the two transmissions complement each other using the Raptor coding. In fact, the reliability enhancement is particularly useful when a user is moving. Nonetheless, the challenge here is how we assign the encoded packets to individual luminaires, so that every received packet from any of luminaire can effectively contribute to decoding.

Suppose there are N_L : $N_L \ll n$ luminaires in a cooperative domain, and they are labeled as $\mathsf{LED}_0, \mathsf{LED}_1, \cdots, \mathsf{LED}_{N_L-1}$. CeilingTalk currently adopts a cyclic assignment scheme in which it assigns all the n encoded packets (codes) to each luminaire but in a circular shift manner, as illustrated in the lower part of Fig. 5. In particular, we set the offset of the i-th luminaire as $o_i = i\lceil n/N_L \rceil, 0 \leq i \leq N_L - 1$, then the sequence numbers of the codes assigned to this luminaire is $\{o_i, o_i + 1, \cdots, n-1, 0, 1, \cdots, o_i - 1\}$. Apparently, this assignment may increase the throughput by N_L times if the reception from each luminaire is perfect, and it certainly allows for quick compensations of lost packets. In a large indoor facility, there can be multiple cooperative domains, and our assignment scheme is applied individually to each of these domains. Since the luminaires are all connected to the same power grid, Power Line Communication (PLC) [28] can be used for scheduling transmissions.

3.3 Pre-Decoding: RoI Extraction

Given a frame captured by a camera, only the banded sections in it contain information to be decoded. Such an information-containing region is termed Region of Interest (RoI) [29]. Due to the **spatial asynchrony** inherent to the transmissions, the location of RoIs within a photo is not known in advance. Though it is visually straightforward to recognize these RoIs, extracting them automatically is far from trivial. Existing proposals [7], [14] resort to sophisticated Computer Vision (CV) techniques, introducing such a high computational overhead that the computations have to be offloaded to a cloudlet server [7].

Our innovative scheme makes use of the default function of the camera to produce RoI masks so that the intensive image processing can be largely avoided. During a continuous reception, the first frame of every p frames is exposed normally (e.g., 1/60s) as shown in Fig. 6(a) and the rest are quickly exposed (e.g., 1/7500s) to obtain Fig. 6(d). The first frame is then converted to a binary RoI mask for extracting the RoIs from the remaining $p - 1$ frames. As a normal exposure would over-expose the luminaires, we convert the first frame to a binary one shown in Fig. 6(b) by setting pixels with Y value (the luma component of YCbCr color space) exceeding threshold Y_t to "1" and otherwise to "0"; here Y_t is normally set to 240 as the Y value for a luminaire under the normal exposure is almost always 255. Then the algorithm builds a rectangular contour for each cluster of "1" pixels to create the RoI mask shown in Fig. 6(c). Obviously, the information to be decoded from the remaining frames can only exist within these contours. As the frames are shot within a very short period of time, we assume there are no tilting/shifting between the mark and the remaining information frames. However, we are also on the way to design an gyroscope-assisted algorithm to compensate for hand motions.

(a) Normally exposed frame (b) Binary frame

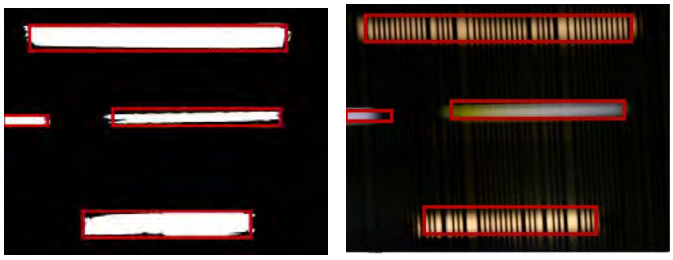

(c) Detected ROIs (d) ROIs for decoding

Fig. 6. Lightweight RoI extraction: for each normally exposed frame (a), we convert it to a binary frame in (b) and detect contours in (c), so that RoIs in a quickly exposed frame (d) can be extracted.

To establish a baseline, we also propose a CV-based RoI extraction scheme. Basically, each frame first goes through a Gaussian blur, and then the aforementioned binary conversion (with a different threshold) and contour detection procedures are used to extract the RoIs. All these functions are carried out by using OpenCV for Android, and this scheme is already a much simplified version of those used in [7], [14]. Obviously, the Gaussian blur procedure is both time and energy consuming.

3.4 Demodulation and Decoding

The traditional demodulation methods often pre-determine a set of sampling times and then compare the samples with a threshold (to distinguish bright and dark bands). Whereas earlier proposals suggest using polynomial regression [8] to determine the threshold due to low SNR, CeilingTalk (with a sufficient level of SNR) achieves a satisfactory detection accuracy by simply setting the threshold as the average of the maximum and minimum pixel illuminance values within an RoI. Nevertheless, the blooming effect of a camera (bleeding or smearing photons from saturated pixels to adjacent pixels) brings difficulty in determining the sampling times. For CeilingTalk, we make demodulation decisions by reasoning on the widths of bright/dark bands. Given the packet structure described in Section 3.1, we can only have four widths for white bands and four for dark bands to be decoded after the preamble, namely $W_{1b}, ..., W_{4b}$ for bright bands and $W_{1d}, ..., W_{4d}$ for dark bands, whereas the preamble has a signature of W_{5b}, where W_{1b} represents the width of one bright band and likewise for other symbols. As the blooming effect diminishes with the number of consecutive bright bands, we derive 5 possible relations between the symbol width W the aforementioned quantities: i) $W_{2b} \leq 2W_{1b}, W_{2d} \geq 2W_{1d}$, ii) $W = W_{1b} + W_{1d}$, iii) $2W = W_{2b} + W_{2d}$, iv) W_{3b} or $W_{3d} \approx 1.5W$, and v) W_{4b} or $W_{4d} \approx 2W$, where $W = W_r F_r / F$ is determined by the

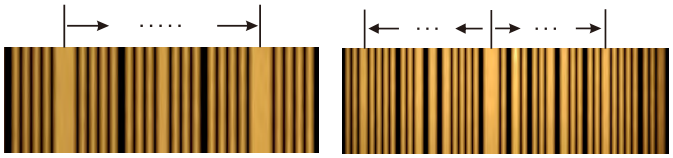

(a) Normal demodulation. (b) CeilingTalk demodulation

Fig. 7. Compare CeilingTalk demodulation with the normal one. (a) Normal demodulation requires two preambles, and (b) CeilingTalk demodulates in for- and back-ward directions starting from the preamble.

camera parameters explained in Section 2.2: F_r is the rolling-shutter frequency, W_r is the rolling-shutter width, and F is the working frequency of an LED luminaire.

Following the aforementioned reasoning, the demodulation algorithm scans through the whole packet to recognize the widths of bands. A normal demodulation scheme identifies and scans a packet as the region between two widest bright bands of width W_{5b} (roughly equal to $2.5W$) as shown in Fig. 7(a). This would, however, sacrifice half of the RoI length (hence halving the throughput) in the worst case given the **temporal asynchrony** in transmissions: a packet could start anywhere within an RoI. Therefore, CeilingTalk employs a bidirectional demodulation process as shown in Fig. 7(b): it starts the scanning from a single widest bright band and collects all the bright and dark bands in both forward and backward directions. All bright/dark bands are clustered into 4 groups, namely $W_{1b}/W_{1d}, \cdots, W_{4b}/W_{4d}$. As there are at most several tens of bands, the clustering can be conducted efficiently. Finally, starting from the widest band again, the bits are recognized in both directions until a predetermined number of bits are decoded. This demodulating flow is shown in Fig. 8. We also give a concrete example

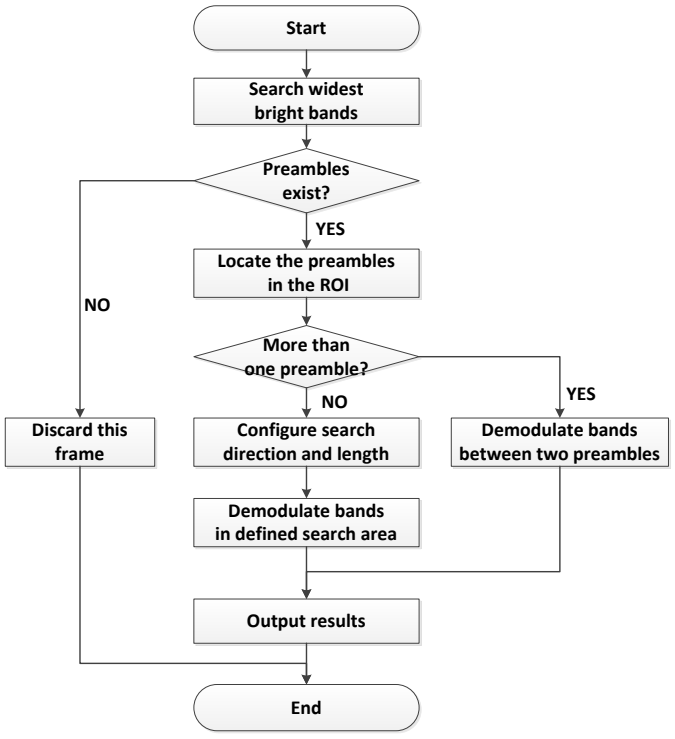

Fig. 8. Demodulation procedure flowing chart.

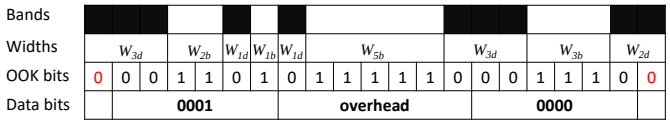

Fig. 9. Demodulation procedure for an 8-bit packet.

in Fig. 9 to demonstrate the demodulation procedure using a packet with 8 bits (12 RLL 4B6B bits) as an example. From the widest bright band, the bits obtained in backward direction are "0001101" and the bits forward are "0011100". As there are 12 RLL bits in total, we can combine the two sequences by removing the duplicate bit "0" and obtain the complete packet as "001110001101" and then "00000001" after RLL 4B6B decoding.

After the demodulation, Raptor decoding procedure takes over with the knowledge of the size of a packet and the number of original packets k. Upon receiving $m = k(1+\epsilon)$ encoded packets ($\epsilon = 0.15$ is chosen to be the overhead of LT codes for CeilingTalk), Gaussian elimination is used for decoding. As one frame may capture multiple RoIs, the decoder waits for all RoIs to be demodulated and combines all received packets in one frame for decoding, so as to reduce the decoding latency dramatically. As mentioned in Sec. 3.2, this allows for almost immediate compensation of packet loss and a manyfold increase in throughput.

3.5 Dimming

Dimming is an important requirement posed by IEEE 802.15.7. As mentioned in 3.1, CeilingTalk can support wide range dimming duty cycle by its hybrid modulation. As the duty cycle of OOK D_{OOK} is fixed given the defined packet structure, we can tune that of PWM D_{PWM} to meet a certain dimming requirement. Specifically, we can tune the overall duty cycle D from 0 to D_{OOK} by embedding the PWM signal into OOK "1":

$$D = D_{\mathrm{OOK}} \times D_{\mathrm{PWM}}, \tag{1}$$

as shown in Fig. 3(a). To reach a full range of D, we can further embed the PWM signal into OOK "0" so that the value of D can go all the way up to 100%:

$$D = D_{\mathrm{OOK}} + (1 - D_{\mathrm{OOK}}) \times D_{\mathrm{PWM}}, \tag{2}$$

as shown in Fig. 3(b). Obviously, making D too close to 0 or 100% can hurt the SNR and thus the throughput, so a preferable range of D for CeilingTalk is from 10% to 70%.

4 Performance Analysis

In this section, we analyze the performance-impacting factors and present the achievable performance of CeilingTalk under realistic constraints.

4.1 Receiver-Camera Configurations

The major issue hampers the performance of CeilingTalk is the inter-symbol interference caused by the blooming effect. Ideally, bright and dark bands should have the same width and the demodulation should be trivial. However, the blooming effect causes a bright band to leak into a neighboring dark band, resulting in irregular widths in bands. While reducing exposure index may suppress this effect, it affects SNR too. So we intend to verify the impacts of various combinations of LED illuminance and camera sensitivity settings on both the SNR (contrast in case of CeilingTalk) and blooming effect.

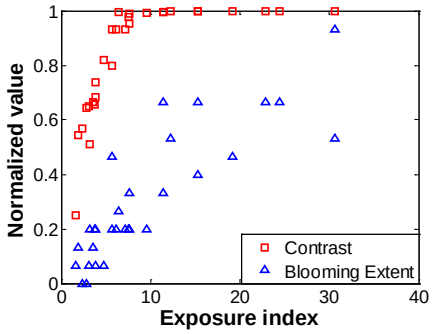

Fig. 10. The image quality with varying exposure indices. A higher exposure index leads to a higher contrast but also a more severe blooming extent.

We have three parameters at hand: illuminance, ISO and exposure time; they control output of each pixel (sensor) by adjusting the incident light intensity, sensor sensitivity, and time to receive photons, respectively. We vary the LED illuminance from 153lux to 301lux (corresponding to those obtained at 1.5m distance with the PWM duty cycle from 10% to 25%), the ISO sensitivity from 100 to 300, and exposure time from $1/7500$s to $1/5000$s, and the exposure index is proportional to the product of illuminance, ISO and the reciprocal of exposure time. In order to unify contrast and blooming effect into the same perspective, we normalize them before plotting the results in Fig. 10. The normalized contract is computed as $(Y_b - Y_d)/Y_{\max}$, with Y_b, Y_d, and $Y_{\max}$ being the luminance values of the bright band, the dark band, and the camera specified maximum. We indicate the blooming effect by W_{1b} and compute its normalized value as $(W_{1b} - 0.5W)/(0.5W)$. Fig. 10 shows that, though the contract increases with the exposure index initially, it quickly gets saturated. At the same time, the blooming effect appears to grows linearly with the exposure index. In fact, when we reach the highest exposure index, W_{1b} reaches up to 32pixels, almost occupying its neighboring dark band. Apparently, the optimal exposure index should be right before the saturation of the contrast, where the blooming effect is still well controlled. Considering that the bands can be even narrower if the LED frequency is further increased to enhance throughput, CeilingTalk could be more sensitive to the blooming effect. Therefore, given a normal indoor illuminance between 200 to 700lux by standard illumination requirement, we tune the PWM duty cycle to meet the requirement, while choosing ISO as 200 and exposure time as $1/7000$s in the office environment where CeilingTalk is deployed and tested; more detailed justifications on these settings will be presented in Sec. 5.2.

4.2 System Model

We hereby explain how we set the packet length packet (in particular the payload length according to the structure specified in Sec. 3.2) and also derive the achievable bit rate given several design parameters.

We first set up the relations between key system parameters and the admissible packet length (thus bit rate), then we deduce their quantities based on the specific devices adopted by our currently implementations. Let us assume i) the LED transmitter has length L long and works at frequency F, ii) the camera receiver has a focal length f_c, sensor width S_w, frame width I_w, column resolution P_c, rolling-shutter frequency F_r and its corresponding width W_r at given column resolution, and frame rate R, and iii) the communication distance is less than d and there are N_L transmitters within this range, we have:

- The projected RoI in a frame has the length of $L_{\text{RoI}} = \frac{\phi L}{d}\sin(a_v) \leq P_c$, where $\phi = f_c I_w / S_w$ and a_v is the view angle.
- The number of OOK bits an RoI can cover should satisfy $N \leq \left\lfloor \frac{L_{\text{RoI}} F}{W_r F_r} \right\rfloor \leq \frac{P_c F}{W_r F_r}$.
- To ensure that an RoI can cover at least one complete packet under asynchronous transmission, the packet size in data bits should satisfy $P_{\text{size}} \leq \lfloor \eta(N-7) \rfloor$, where η is the encoding efficiency of employed modulation scheme, e.g. $\eta = 4/6$ for 4B6B, and each packet has 7 bits of overhead.
- According to above analysis, we can bound the bit rate of the CeilingTalk as:

$$\begin{aligned} C = N_L R P_{\text{size}} &\leq N_L R\Big(\left\lfloor \eta \times \frac{L_{\text{RoI}} F - 7}{W_r F_r} \right\rfloor \Big) \\ &\leq N_L R\Big(\left\lfloor \eta \times \frac{P_c F - 7}{W_r F_r} \right\rfloor \Big) \end{aligned} \tag{3}$$

Now we can put these parameters into practical perspectives in order to check the actual quantities of packet size and achievable bit rate of CeilingTalk. Given the rear-facing camera of Nexus 6 as the CeilingTalk receiver, we have $R = 30$fps and $F_r = 19$kHz as we measured, $W_r = 3$pixel (i.e., three columns are exposed at one time) under a preview resolution of 1920×1440, which allows us to derive $\phi = 2325$pixel. For the front camera, the parameters are slightly different from the rear camera: $W_r = 4$pixel and hence $\phi = 2007$pixel under preview resolution of 1024×768. As the working frequency of LED can go up to a few GHz, the working frequency of CeilingTalk is restricted mainly by the capability of the receiver, namely the rolling-shutter frequency F_r of the camera.

Since we measured $F_r = 19$kHz, we have to set $F \leq F_r/2 = 9.5$kHz. We choose $F = 8$kHz for reliability purpose, which results in a symbol width $W = 11$pixel at preview resolution of the rear camera and $W = 9$pixel at preview resolution of the front camera. Apparently, the working frequency is set to meet the the decodability at the preview resolution, as we need some "guard pixels" to combat the blooming effect. In an indoor environment, the ceiling is around 3m from a hand-held smartphone camera sitting right below it. Considering a 0 to 4m horizontal distance, we have $d_{rc} = 5$m as our maximum communication distance for rear camera. As for the front camera, normally we hold the phone in parallel with the floor, thus we regulate $d_{fc} = 3$m as the maximum communication distance for front camera.

Substituting $F = 8$kHz and $d_{rc} = 5$m and $d_{fc} = 3$m into earlier formulas, we obtain suitable $N \leq 78$ for both rear and front cameras and thus maximum $P_{\text{size}} = 44$bit.As we need a few guard bits to combat tilting of the phone or RoI extraction inaccuracy, we set the actual $P_{\text{size}} = 40$bit. Therefore, CeilingTalk can offer a bit rate of 1200bps with one transmitter at the maximum communication distance, given the frame rate $R = 30$fps of the cameras, as summarized in TABLE 3 [1]. This bit rate can be further improved by combining more transmitters and using the full resolution of the camera. For example, if we have 6 transmitters within the communication range and we use either front camera or rear camera for decoding, the bit rate can be boosted to 7200bps in theory.

For a single transmitter, there are other ways to improve the bit rate. According to Eq. (3), reducing the communication distance can increase the bit rate in proportion. Also, using a smartphone camera with a higher resolution and/or frame rate can also proportionally increase the bit rate. For example, if we use iPhone 5s with a frame rate 120fps, the bit rate can be tripled.

TABLE 3
System setting and achievable capacity

	Rear-facing	Front-facing
Resolution	1920 × 1440	1024 × 768
F	8kHz	8kHz
W	11pixel	9pixel
d	5m	3m
N	78	84
P_{size}	40	40
C	1200bps	1200bps

4.3 Practical Considerations

The above discussions assume rather ideal scenarios, but various interferences exist in a realistic deployment. One may expect the ambient lighting to be a major interference to CeilingTalk communications, but the fact is that ambient lighting only causes noticeable noise if a light sensor is used as the receiver [10] or if the rolling-shutter sensing is applied to some surface reflection [8], [9]. In our case, as CeilingTalk directly uses the LED luminaires as the signal source, the SNR is so high that we can run a receiver with an extremely short exposure time, suppressing the interference (ambient light "leaking into" the RoIs) to the largest extent. Therefore, the interference from ambient lighting can surely be neglected for CeilingTalk.

Although CeilingTalk regulates each packet to be transmitted within periodic 33.3ms time frames (i.e. the video frame length captured by the camera), missing packets are inevitable due to varying inter-frame gaps and the gap jitter induced transmitter-receiver asynchrony. Raptor codes shows its advantage again in combating the asynchrony: since the packet loss caused by inter-frame gaps (and their jitters) is about 5.6% as we measured, the overhead 25% that we use for Raptor encoding is sufficient to generate repair packets for compensating the packet loss caused by varying inter-frame gaps.

1. Although the current prototype of CeilingTalk is using Nexus 6 as the receiver, it can be readily extended to diversified smartphones and the parameters are set by the device with the lowest performance.

(a) LED luminaire made of commercial LED strips

(b) VLC control board and LED driver

Fig. 11. CeilingTalk transmitter.

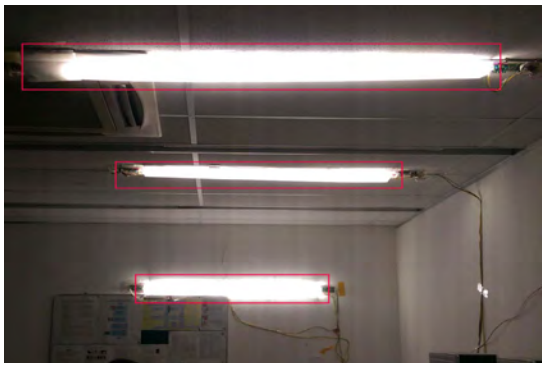

Fig. 12. CeilingTalk testbed: three LED luminaires are mounted on wall and ceiling with an interval of 1.5m, and the ceiling is 2.5m above the floor.

One major challenge CeilingTalk faces is user motion, in particular hand micro-motion caused by user mobility during data reception. As we have cooperative transmissions to handle the user mobility during reception, we shall only focus on the interference caused by hand micro-motion. Note that if a user totally moves out of a cooperative domain, the reception will fail but this is indeed the purpose of CeilingTalk's location-bound communication: the data service is location dependent.

For the hand micro-motion, we assume that its magnitude so minor that it does not cause mis-capturing of the transmitters, which is reasonable for users with normal physical conditions. Therefore, the major problem caused by hand micro-motion is twofold: on one hand, minor camera titling/rotating may distort the RoIs extracted from the first normal exposed frame, causing bit loss within an RoI. On the other hand, major camera tilting may cause certain RoIs un-decodable. While the minor changes can be compensated by gyroscope-assisted geometric transformations, we certainly need to constrain the major ones. Given a tilting angle α, we obtain the length of RoI $\frac{\phi L}{d} \times \cos(\alpha)$. Consequently, the bit rate in Eq. 3 is constrained by $C \leq \cos(\alpha) N_L R \left\lfloor \eta \times \frac{\phi L F - 7}{d W_r F_r} \right\rfloor$. In our current implementation, we choose $P_{\text{size}} = 40$bit that is 60 OOK bits subject to $N = 78$ OOK bits, so the rotation angle should be less than $\arccos(60/78) \approx 40°$ without incurring bit loss. Fortunately, we can let the application to alert an user for such packet loos and hence rely on user assistance to improve the performance.

5 Evaluation

We report the evaluation results on our CeilingTalk testbed in this section. We first explain the experiment settings, and then present the performance evaluations of CeilingTalk with respect to various parameters.

5.1 Experiment Settings

We build CeilingTalk transmitters with commercial LED strips [30] and self-developed LED driver. Fig. 11 shows one such transmitter; it includes an LED luminaire, an driver board, and a control unit. The LED luminaire is made of 16 LED strips each containing 36 LED chips, hence a dimension of 120cm×8cm similar to a common fluorescent luminaire. As this LED luminaire has a nominal drive current of 600mA, we use low-cost transistors to build the driver circuit. To enhance the stability and readability of the transmitter, optocouplers are employed between the control unit and driver.

In total we deploy 3 such transmitters every 1.5m in parallel in our lab, with one on the wall and the other two on the ceiling, as shown in Fig. 12. The one on the wall is deployed for performance evaluation at different transmission distances, while the other two are meant to emulate a realistic setting. We use a Nexus 6 smartphone as the receiver, with parameters specified in Sec. 4.2.

A message is divided to blocks, each contains k packets. Raptor code [27] encodes at the rate of $n = 1.25 \times k$. Given that the PSN field in a packet contains up to 8 bits, each block can have at most 256 coded packets. The decoding overhead is set $\epsilon = 0.15$. Each packet is transmitted repeatedly within a frame duration 33.3ms.

For the dimming level setting, we vary the duty-cycle of PWM within its full range and ask 15 volunteers to observe the LED luminaires for dimming changes and flicker. We find that CeilingTalk's OOK-PWM modulation scheme can support a full range duty-cycle control, but when the duty-cycle is less than 3% , i.e. at very low dimming levels, the flicker becomes observable. Given that very low dimming levels are rare in the indoor environments, we fix the duty cycle of LED luminaires hereafter at 20% (roughly 250lux at 1.5m distance, a level with which all volunteers feel comfortable) unless otherwise stated.

5.2 Camera Settings

In Sec. 4 we have analyzed how to choose the exposure time and ISO of the camera to strike a balance between high contrast and low blooming effect. Here we focus on the impact of camera setting on Bit Error Rate (BER) and Packet Error Rate (PER); both of them take the missing frames into account. In Fig. 13, we use Root Mean Square Error (RMSE) to measure the fluctuation in the widths of the bands. It indicates the difference between the widths of all detected bands and the widths of bands leading to accurate bit decision. RMSE has direct impact on BER because the width of a band serves as the criterion for decoding. In Fig. 13(a), we fix the exposure time at $1/7000$s, the LED duty cycle at 20%, and communication distance at $3m$. With ISO varying from 100 to 600, RMSE and BER first decrease as the contrast

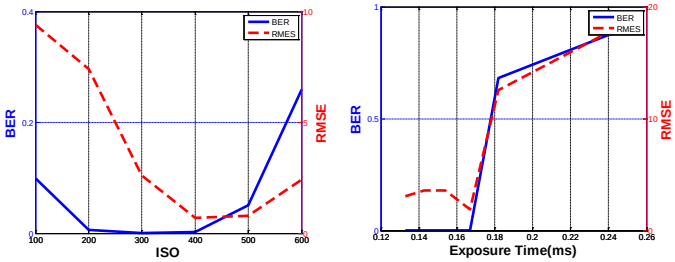

(a) BER vs. ISO (b) BER vs. Exposure time

Fig. 13. The impact of camera settings. (a) BER reaches the lowest at ISO from 200 to 400. (b) BER increases with the exposure time.

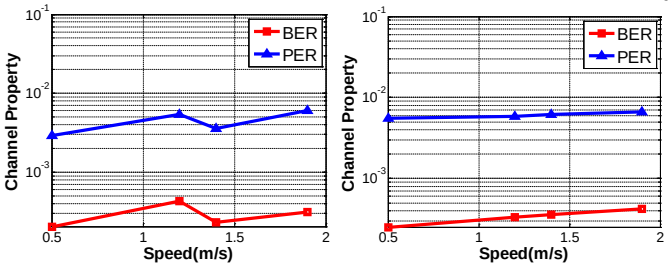

(a) CV-based decoding with rear camera (b) CV-based decoding with front camera

Fig. 15. Channel properties under mobile scenarios. Both BER and PER stay at around 10^{-3} when users moving within a reasonable communication range.

is enhanced and then degrade as the blooming effect causes more irregular widths of bands. Fig. 13(b) illustrates how the exposure time affects the performance. We fix the ISO at 200 and other parameters as before. Due to the same reason of blooming effect, RMSE and BER remain stably low at short enough exposure time and then increase when the exposure time exceeds $1/6000s$. The result confirms to our analysis in Section 4. Therefore, CeilingTalk by default sets ISO at 200 and exposure time at $1/7000s$ unless otherwise specified.

5.3 Channel Property

We first evaluate the channel properties with the camera fixed at a stationary position. In Fig. 14 we measure the channel properties in terms of BER and PER at different communication distances from 1.5m to 5.5m for rear camera and from 1m to 3.5m for front camera, and we also compare the performance of our Non-CV-based decoding method with the baseline CV-based decoding. We observe that both BER and PER stay at a low level regardless of the increasing communication distance until the $d > 5$m for rear camera and $d > 3$m for front camera, respectively. At $d = 5.5$m and $d = 3.5$m, BER and PER of both rear and front cameras dramatically increase because the captured RoI becomes too small to cover P_{size}, thus causing non-negligible decoding failures. Moreover, the performance in terms of BER and PER of the two decoding methods appear to be very similar, proving the superiority of our Non-CV-based method: the computation time is about 9ms in our case, whereas the baseline CV-based method takes about 30ms to complete.

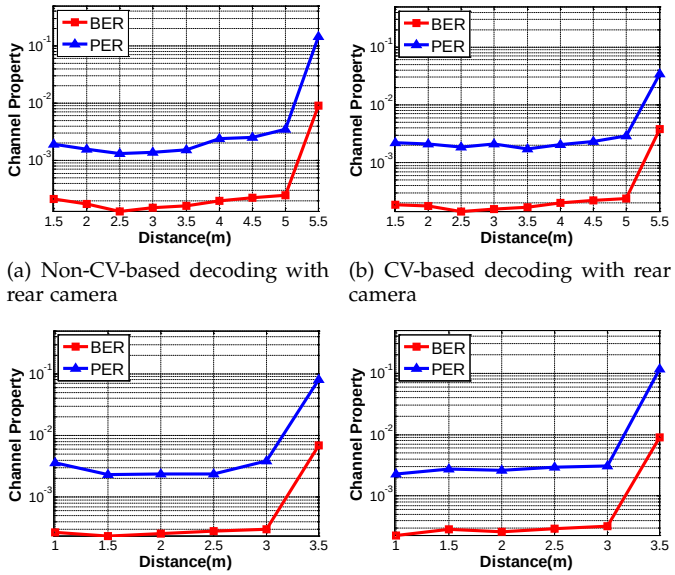

(a) Non-CV-based decoding with rear camera (b) CV-based decoding with rear camera

(c) Non-CV-based decoding with front camera (d) CV-based decoding with front camera

Fig. 14. Channel properties under stationary scenarios. The channel is stable until the communication distance goes beyond 5m and 3m for rear and front cameras, respectively.

We then evaluate the channel properties under mobile scenarios. We only use the CV-based decoding here due to drastic changes in the camera position. This experiment is conducted by two users. They hold the smartphone and walk arbitrarily within the communication range with different speed, while guaranteeing that the camera can always capture the targeted LED luminaire. During the experiment, the accelerometer on the phone is turned on to monitor the average walking speed. Fig. 15 shows that the BER and PER under mobile scenarios are comparable to (albeit slightly lower than) those of stationary scenarios, confirming the robustness of CeilingTalk against mobility. Both these two experiments allow us to conclude that CeilingTalk provides stable channels with very low BER/PER ($\sim 10^{-3}$) within a reasonable communication range for both rear and front cameras. Note that the non-monotonic trends of BER/PERs shown in both Fig. 14 and Fig 15 can be attributed to the non-uniformity in performing different sets of experiments. Therefore, we should interpret the results as roughly constant BER/PERs within a reasonable d.

5.4 Throughput and Latency

In this section, we evaluate the benefit of using Raptor codes in terms of throughput and latency, by comparing CeilingTalk with a baseline Non-Encoding mechanism that simply sends the set of k original packets assigned to individual lumiaires in a circular shift manner. The latency is the elapsed time from when the camera starts to receive data to when it recovers (for CeilingTalk) or receives (for the baseline mechanism) all original packets. In an LED-camera communication session, the latency is determined by the frame rate of the camera since the frames are taken about every 33.3ms that dominates other factors such as decoding computations. Therefore, the latency is approximately proportional to the numbers of camera frames needed for successful decoding. The throughput is computed as the total bits of the k packets divided by the latency. We conduct

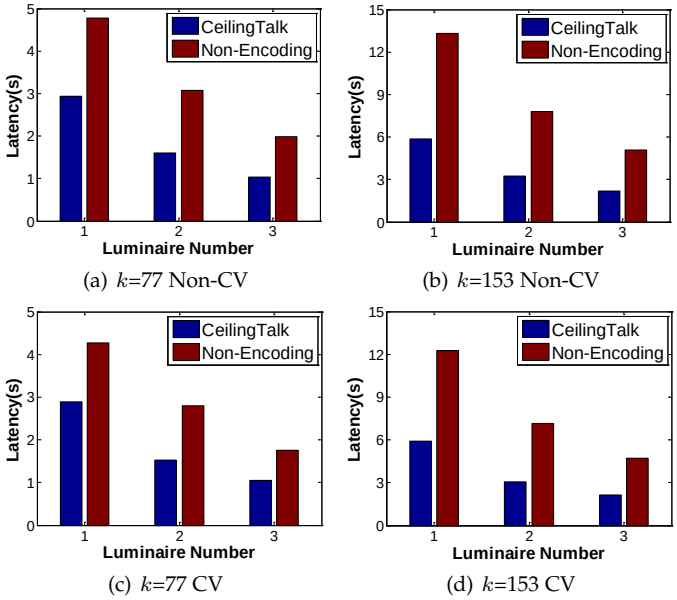

(a) k=77 Non-CV (b) k=153 Non-CV

(c) k=77 CV (d) k=153 CV

Fig. 16. Latency vs. luminaire number with front camera. The improvement in terms of latency grows with the luminaire number.

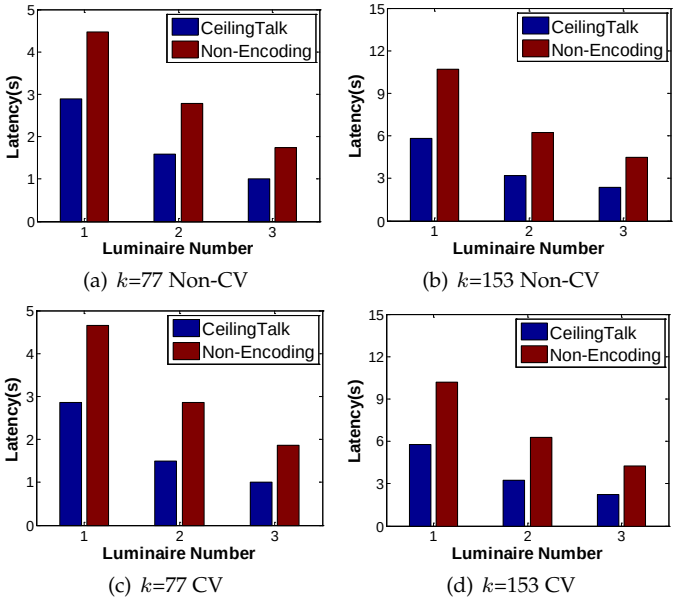

(a) k=77 Non-CV (b) k=153 Non-CV

(c) k=77 CV (d) k=153 CV

Fig. 17. Latency vs. luminaire number with rear camera. The improvement in terms of latency grows with the luminaire number.

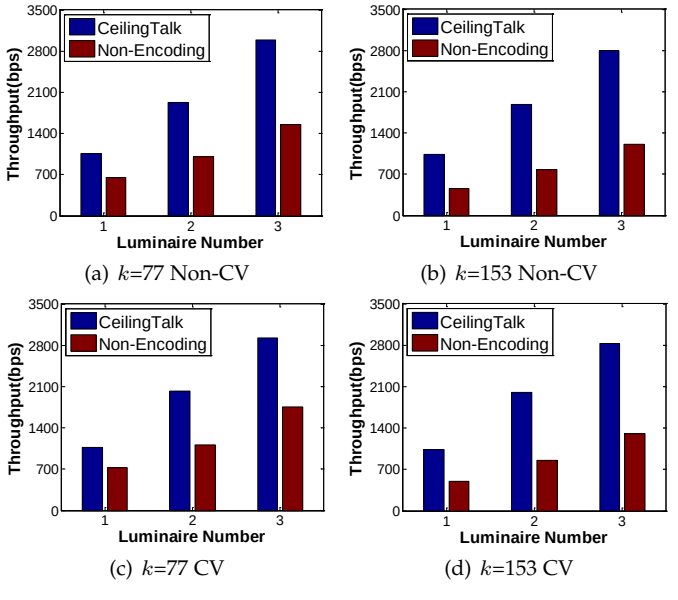

(a) k=77 Non-CV (b) k=153 Non-CV

(c) k=77 CV (d) k=153 CV

Fig. 18. Throughput vs. luminaire number with front camera. The improvement in terms of throughput grows with the increasing luminaire number.

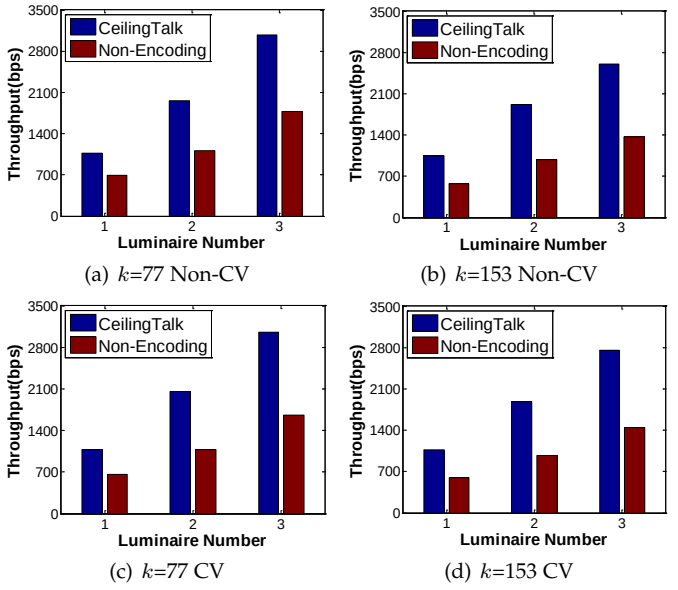

(a) k=77 Non-CV (b) k=153 Non-CV

(c) k=77 CV (d) k=153 CV

Fig. 19. Throughput vs. luminaire number with rear camera. The improvement in terms of throughput grows with the increasing luminaire number.

the experiments with $k = 77$ and $k = 153$ to evaluate how CeilingTalk performs under different k values.

We first evaluate the performance of CeilingTalk under stationary scenarios as explained in Sec. 5.3. The achieved latency by front and rear cameras is plotted in Fig. 16 and Fig. 17, respectively. The achieved throughput by front and rear cameras is plotted in Fig. 18 and Fig. 19. We can observe that both metrics get improved proportionally to the number of luminaires, and the throughput of both front and rear cameras can reach 3.0kbps roughly with 3 luminaires. In fact, the rear camera of CeilingTalk has the potential to capture at most 6 luminaires in one frame, so it can offer a throughput beyond 6.0kbps.

We can also observe that, CeilingTalk brings more significant improvement when $k = 153$ than that when $k = 77$ compared with the baseline mechanism. As shown in Fig. 18 and Fig. 19, the gain of CeilingTalk over Non-Encoding mechanism rises from about 50% under $k = 77$ to about 80% under $k = 153$ with one luminaire and 50% to almost 100% with three luminaires. Therefore, the improvement also grows with k. This significant improvement stems from the increased packet loss rate under a larger k, for which the benefit of Raptor coding becomes more evident.

We also perform evaluations under mobile scenarios where we deliberately let the phone cameras of two users (A and B) mis-capture one or two luminaires from time to time, and the results are reported in TABLE 4 and TABLE 5.

TABLE 4
Performance under mobile scenarios using front camera

	CeilingTalk		Non-Encoding	
k	Latency	Throughput	Latency	Throughput
77 (A)	1.68s	1.84kbps	4.67s	0.66kbps
77 (B)	1.77s	1.75kbps	5.21s	0.59kbps
153 (A)	3.89s	1.57kbps	17.74s	0.35kbps
153 (B)	4.13s	1.48kbps	15.88s	0.39kbps

TABLE 5
Performance under mobile scenarios using rear camera

	CeilingTalk		Non-Encoding	
k	Latency	Throughput	Latency	Throughput
77 (A)	1.50s	2.05kbps	3.73s	0.83kbps
77 (B)	1.58s	1.96kbps	4.21s	0.73kbps
153 (A)	3.43s	1.78kbps	13.38s	0.46kbps
153 (B)	3.34s	1.83kbps	12.37s	0.49kbps

We have almost the same observations as those under stationary scenarios that CeilingTalk outperforms the baseline mechanism and the benefit grows with increasing k. In TABLE 5 user A with rear-facing camera and $k = 77$ obtains an average throughput 2.05kbps with CeilingTalk but only 0.83kbps otherwise, while user B obtains 1.96kbps against 0.73kbps otherwise. TABLE 4 shows the similar performance by using front camera. Therefore, these results again confirm the benefit of employing Raptor codes for enhancing reliability.

5.5 Power Consumption

The power consumption of CeilingTalk has two parts: the consumption of LED control board and that of camera for frame capture and decoding. Since CeilingTalk utilizes an existing lighting infrastructure as transmitters, the consumption of the communication front-end (the luminaires) is actually zero. The real consumption is caused by receiving data from Ethernet or PLC backbone, as well as the encoding computations. Our field experiments show that the driver of CeilingTalk has a power consumption of 40mW, but this consumption appears to be constant regardless of encoding computations, suggesting that it is mainly the consumption of driving light emission (the default function of a luminaire). Therefore, we can conclude that the tx power consumption of CeilingTalk is negligible. As for Wi-Fi, existing Wi-Fi routers consume at least 800mW according to [31], [32].

The receiver of CeilingTalk does lead to a rather high power consumption due to the use of image sensors. According to [33], the image sensor undergoes two states within a frame duration T_{frame}, i.e. active mode and idle mode. Given the duration and power for the two modes per frame, i.e. T_{active}, P_{active}, T_{idle} and P_{idle} respectively, the power of CeilingTalk is formulated as $P = (P_{active}T_{active} + P_{idle}T_{idle})/T_{frame}$. In our case, based on our understanding of the information provided by Sony, we calculate $P_{active} \approx 183$mW, $T_{active} = N/f$ where N is the resolution and $f > 338$MHz is pixel rate, $P_{idle} \approx 77$mW, $T_{idle} = T_{frame} - T_{active}$ and $T_{frame} = 33$ms. With $N = 1024 \times 768$ used by CV-based RoI extraction and $N = 1920 \times 1440$ used by lightweight RoI extraction, we obtain the power $P < 84$mW and $P < 103$mW respectively. However, by using aggressive standby mechanism which lets the image sensor work in standby mode when no operation is performed [33], T_{idle} goes down to $T_{exposre} = 1/7000$s so that P drops to $P \approx 20$mW and 51mW respectively. Moreover if optimal clock scaling is adopted, the power can be reduced further. Wi-Fi appears to be more efficient at the receiver side by consuming slightly more than 50mW for data reception [34], so the receiver side energy consumption is a major limitation of CeilingTalk. We summarize these quantities in Table 6. In fact, as a data service, the transmitter side may constantly consume energy while the receiver side only causes consumption intermittently. Therefore, CeilingTalk is much more energy efficient than Wi-Fi in terms of the overall system consumption.

TABLE 6
CeilingTalk vs. Wi-Fi: Comparison of Power Consumption.

	Transmitter (mW)	Receiver (mW)	Total (mW)
CeilingTalk	~ 0	<103 [33]	<103
Wi-Fi	>800 [32]	50 [34]	>850

6 CONCLUSION

In this paper, we have presented CeilingTalk as an LED-camera VLC system. It innovates in both encoding and decoding schemes to improve link reliability and throughput, so that it allows us to have the first realistic LED-camera VLC deployment, and it also provides us with practical insights on how such systems should be configured to reach its maximum capacity. Extensive field experiments have shown that our system can achieve a throughput much higher than a recent experimental prototype [14]. Our future work aims to further improve CeilingTalk in terms of throughput and energy efficiency, by designing more effective coding/decoding schemes.

REFERENCES

[1] J. Hao, Y. Yang, and J. Luo, "CeilingCast: Energy Efficient and Location-Bound Broadcast Through LED-Camera Communication," in *Proc. of the 35th IEEE INFOCOM*, 2016, pp. 1629–1637.

[2] S. Rajagopal, R. Roberts, and S. Lim, "IEEE 802.15.7 Visible Light Communication: Modulation Schemes and Dimming Support," *IEEE Communications Magazine*, vol. 50, no. 3, pp. 72–82, 2012.

[3] R. Drost and B. Sadler, "Constellation Design for Channel Precompensation in Multi-Wavelength Visible Light Communications," *IEEE Trans. on Antennas and Propagation*, vol. 62, no. 6, pp. 1995–2005, 2014.

[4] H. Burchardt, N. Serafimovski, D. Tsonev, S. Videv, and H. Hass, "VLC: Beyond Point-to-Point Communication," *IEEE Communications Magazine*, vol. 52, no. 7, pp. 98–105, 2014.

[5] A. Ganick and D. Ryan, "ByteLight Scalable Indoor Location," http://www.bytelight.com.

[6] N. Rajagopal, P. Lazik, and A. Rowe, "Visual Light Landmarks for Mobile Devices," in *Proc. of the 13th ACM IPSN*, 2014, pp. 249–260.

[7] Y. Kuo, P. Pannuto, K. Hsiao, and P. Dutta, "Luxapose: Indoor Positioning with Mobile Phones and Visible Light," in *Proc. of the 20th ACM MobiCom*, 2014, pp. 447–458.

[8] C. Danakis, M. Afgani, G. Povey, I. Underwood, and H. Haas, "Using a CMOS Camera Sensor for Visible Light Communication," in *Proc. of 31th IEEE GLOBECOM*, 2012, pp. 1244–1248.

[9] N. Rajagopal, P. Lazik, and A. Rowe, "Hybrid Visible Light Communication for Cameras and Low-Power Embedded Devices," in *Proc. of 1st ACM VLCS*, 2014, pp. 33–38.

[10] L. Li, P. Hu, C. Peng, G. Shen, and F. Zhao, "Epsilon: A Visible Light Based Positioning System," in *Proc. of the 11th ACM NSDI*, 2014, pp. 331–343.

[11] S. Rajbhandari, H. Chun, G. Faulkner, K. Cameron, A. V. Jalajakumari, R. Henderson, D. Tsonev, M. Ijaz, Z. Chen, H. Haas *et al.*, "High-Speed Integrated Visible Light Communication System: Device Constraints and Design Considerations," *IEEE Journal on Selected Areas in Communications*, vol. 33, no. 9, pp. 1750–1757, 2015.

[12] M. S. Mossaad, S. Hranilovic, and L. Lampe, "Visible Light Communications using OFDM and Multiple LEDs," *IEEE Trans. on Communications*, vol. 63, no. 11, pp. 4304–4313, 2015.

[13] J. C. Chau, C. Morales, and T. D. Little, "Using Spatial Light Modulators in MIMO Visible Light Communication Receivers to Dynamically Control the Optical Channel," in *Proc. of the 13th EWSN*, 2016, pp. 347–352.

[14] H. Lee, H. Lin, Y. Wei, H. Wu, H. Tsai, and K. Lin, "RollingLight: Enabling Line-of-Sight Light-to-Camera Communications," in *Proc. of the 13th ACM MobiSys*, 2015, pp. 167–180.

[15] P. Hu, P. H. Pathak, X. Feng, H. Fu, and P. Mohapatra, "ColorBars: Increasing Data Rate of LED-to-Camera Communication using Color Shift Keying," *Proc. of the ACM CoNEXT*, pp. 1–12, 2015.

[16] B. Zhang, K. Ren, G. Xing, X. Fu, and C. Wang, "SBVLC: Secure Barcode-Based Visible Light Communication for Smartphones," *IEEE Trans. on Mobile Computing*, vol. 15, no. 2, pp. 432–446, 2016.

[17] T. Hao, R. Zhou, and G. Xing, "COBRA: Color Barcode Streaming for Smartphone Systems," in *Proc. of the 10th ACM MobiSys*, 2012, pp. 85–98.

[18] S. Perli, N. Ahmed, and D. Katabi, "PixNet: Interference-Free Wireless Links using LCD-Camera Pairs," in *Proc. of the 16th ACM MobiCom*, 2010, pp. 137–148.

[19] W. Du, J. Liando, and M. Li, "SoftLight: Adaptive Visible Light Communication over Screen-Camera Links," in *Proc. of the 35th IEEE INFOCOM*, 2016, pp. 1620–1628.

[20] A. Wang, S. Ma, C. Hu, J. Huai, C. Peng, and G. Shen, "Enhancing Reliability to Boost the Throughput over Screen-Camera Links," in *Proc. of the 20th ACM MobiCom*, 2014, pp. 41–52.

[21] W. Hu, J. Mao, Z. Huang, Y. Xue, J. She, K. Bian, and G. Shen, "Strata: Layered Coding for Scalable Visual Communication," in *Proc. of the 20th ACM MobiCom*, 2014, pp. 79–90.

[22] W. Hu, H. Gu, and Q. Pu, "LightSync: Unsynchronized Visual Communication Over Screen-Camera Links," in *Proc. of the 19th ACM MobiCom*, 2013, pp. 15–26.

[23] W. Huang and W. Mow, "PiCode: 2D Barcode with Embedded Picture and ViCode," in *Proc. of the 19th ACM MobiCom*, 2013, pp. 139–242.

[24] T. Li, C. An, X. Xiao, A. Campbell, and X. Zhou, "Real-Time Screen-Camera Communication Behind Any Scene," in *Proc. of the 13th ACM MobiSys*, 2015, pp. 197–211.

[25] A. V. Oppenheim, R. W. Schafer, J. R. .Buck *et al.*, *Discrete-Time Signal Processing*. Prentice hall Englewood Cliffs, NJ, 1989, vol. 2.

[26] A. Shokrollahi, "Raptor Codes," *IEEE Trans. on Information Theory*, vol. 52, no. 6, pp. 2551–2567, 2006.

[27] M. Luby, A. Shokrollahi, and T. Stockhammer, "The Open Source Implementation of Raptor Code RFC5053," https://code.google.com/p/raptor-code-rfc/.

[28] H. Ma, L. Lampe, and S. Hranilovic, "Integration of Indoor Visible Light and Power Line Communication Systems," in *Proc. of the 17th IEEE ISPLC*, 2013, pp. 291–296.

[29] G. Corbellini, K. Aksit, S. Schmid, S. Mangold, and T. Gross, "Connecting Networks of Toys and Smartphones with Visible Light Communication," *IEEE Communications Magazine*, vol. 52, no. 7, pp. 72–78, 2014.

[30] "LUXEON XF-3535L," http://www.lumileds.com/products/matrix-platform/luxeon-xf-3535l.

[31] "WiFi Router Power Consumption Database," http://www.tpcdb.com/list.php?type=11.

[32] G. Palem and S. Tozlu, "On Energy Consumption of Wi-Fi Access Points," in *Proc. of the 7th IEEE CCNC*, 2012, pp. 434–438.

[33] R. LiKamWa, B. Priyantha, M. Philipose, L. Zhong, and P. Bahl, "Energy Characterization and Optimization of Image Sensing toward Continuous Mobile Vision," in *Proc. of the 11th ACM MobiSys*, 2013, pp. 62–82.

[34] A. Carroll and C. Heiser, "An Analysis of Power Consumption in a Smartphone," in *USENIX Annual Technical Conference*, vol. 14, 2010.

Yanbing Yang received his BS degree from University of Electronic Science and Technology of China, China, in 2010, and the MS degree from University of Electronic Science and Technology of China, China, in 2013. He is currently a PhD student in the School of Computer Science and Engineering, Nanyang Technological University, Singapore. His research interests are visible light communication, visible light sensing, as well as their applications.

Jie Hao received her BS degree from Beijing University of Posts and Telecommunications, China, in 2007, and the PhD degree from University of Chinese Academy of Sciences, China, in 2014. From 2014 to 2015, she has worked as post-doctoral research fellow in the School of Computer Engineering, Nanyang Technological University, Singapore. She is currently an Assistant Professor at College of Computer Science and Technology, Nanjing University of Aeronautics and Astronautics, China. Her research interests are wireless sensing, visible light communication and etc.

Jun Luo received his BS and MS degrees in Electrical Engineering from Tsinghua University, China, and the PhD degree in Computer Science from EPFL (Swiss Federal Institute of Technology in Lausanne), Lausanne, Switzerland. From 2006 to 2008, he has worked as a post-doctoral research fellow in the Department of Electrical and Computer Engineering, University of Waterloo, Waterloo, Canada. In 2008, he joined the faculty of the School of Computer Science and Engineering, Nanyang Technological University in Singapore, where he is currently an Associate Professor. His research interests include mobile and pervasive computing, wireless networking, applied operations research, as well as network security. More information can be found at http://www.ntu.edu.sg/home/junluo.