

Communication Networks of Visible Light Emitting Diodes with Intra-Frame Bidirectional Transmission

Qing Wang
IMDEA Networks Institute & UC3M
Madrid, Spain
qing.wang@imdea.org

Domenico Giustiniano
IMDEA Networks Institute
Madrid, Spain
domenico.giustiniano@imdea.org

ABSTRACT

Unlike traditional radio frequency communication of consumer devices, the “optical antenna” direction of Visible Light Communication (VLC), *i.e.*, the Field-Of-View (FOV), varies greatly from device to device. This encompasses wide FOVs of ambient infrastructure and directional FOVs of light emitted by low-end embedded devices. This variety of light wave propagation can severely affect the transmission reliability, despite “pointing” devices to each other may seem enough for a reliable link. In particular, the fact that FOVs are unknown makes traditional access protocols in VLC unreliable in presence of interference among nodes of different FOVs and exacerbates the hidden-node problem. In this paper, we propose a Carrier Sensing Multiple Access/Collision Detection&Hidden Avoidance (CSMA/CD-HA) Medium Access Control protocol for a network where each node solely uses one Light Emitting Diode (LED) to transmit and receive data. The CSMA/CD-HA can enable *in-band intra-frame* bidirectional transmission with just one optical antenna. The key idea is to exploit the intra-frame data symbols without emission of light to introduce an embedded communication channel. This approach enables the *transmission* of additional data *while receiving* in the same optical channel and it makes the communication robust to different types of FOVs. We build a software-defined embedded platform running on Linux operating system, implement the CSMA/CD-HA protocol, and evaluate its performance through experiments. Results show that collisions caused by hidden nodes can be reduced and our protocol can increase the saturation throughput by nearly up to 50% and 100% under the two-node and four-node scenarios, respectively.

Categories and Subject Descriptors

C.2.1 [Computer-Communication Networks]: Network Architecture and Design—*Wireless Communication*

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

CoNEXT'14, December 2–5, 2014, Sydney, Australia.

Copyright 2014 ACM 978-1-4503-3279-8/14/12 ...\$15.00.

<http://dx.doi.org/10.1145/2674005.2675000>.

Keywords

Visible light communication networks; CSMA/CD-HA; Bidirectional transmission; Design; Implementation; Evaluation

1. INTRODUCTION

Visible Light Communication (VLC) is emerging as a complementary technology to mainstream research on Radio Frequency (RF) communication. VLC utilizes visible light from Light Emitting Diodes (LEDs) to convey digital information between devices. A network of visible LEDs could be enabled by connecting various devices such as ceiling bulbs, lamps, light emitters embedded into cars and mobile devices and perhaps, in the future, LED TVs. However, the directionality of the “optical antenna” of VLC, *i.e.*, the Field Of View (FOV) of LEDs, varies greatly from device to device. An infrastructure, *e.g.*, a light bulb on the ceiling, normally emits light with a wide-FOV. In contrast, mobile devices may have various FOVs, according to the space and power constraints.

This variety of light wave propagation calls for networking approaches that are robust to the specific optical antenna. The design opportunities can take advantage of two fundamental differences with respect to RF communication.

- First, VLC often adopts the On-Off Keying (OOK) modulation, thus a transmitter can be “idle” when transmitting an “OFF” signal. This implies that other communications could be established during these short times without light emission. These concurrent communications may improve the system performance, such as increasing the throughput and so on.
- Second, while photodiodes are normally used as receivers, a LED has been proved to work as a receiver in LED-to-LED communications [5, 6]. Thus, a network of LEDs would only require one LED as optical antenna at each transceiver. The challenge is then how to create a network of LEDs *with* different FOVs and *without* additional optical components.

As illustrated in Fig. 1, these differences *i)* bring opportunities to design new Medium Access Control (MAC) protocols to improve the network performance; *ii)* increase the hidden-node problem in some scenarios. In this paper, we propose a Carrier Sensing Multiple Access/Collision Detection&Hidden Avoidance (CSMA/CD-HA) Medium Access Control (MAC) protocol to enable intra-frame bidirectional VLC in a network of LEDs. Similarly to the emerging full-duplex of RF communications with a single antenna proposed by [3], CSMA/CD-HA only uses one “antenna” – a

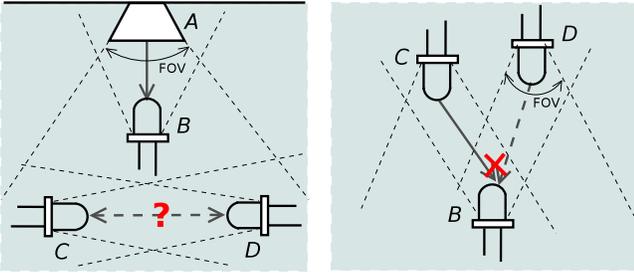


Figure 1: Motivation – various “optical antenna” directions of VLC nodes and OOK modulation: *left*) is in-band interference-free concurrent transmissions of $A \rightarrow B$ and $C \leftrightarrow D$ possible? *right*) exacerbate the hidden-node problem.

single LED to enable a secondary channel. In contrast to full-duplex research in RF domain, we target visible light spectrum and implement very simple techniques to enable the secondary channel with off-the-shelf embedded hardware that runs on credit-card-sized embedded board with a total unit cost of approximately 50 dollars. The proposed protocol can also alleviate the hidden-node problem without using the Request-To-Send/Clear-To-Send (RTS/CTS) method. To evaluate the performance of our protocol, we design a software-defined embedded platform for visible light communication networks, which runs Linux and implements CSMA/CD-HA as part of it. Experimental results show that CSMA/CD-HA can increase the saturation throughput by nearly up to 50% and 100% under the two-node and four-node scenarios, respectively, with respect to state-of-the-art MAC protocols.

2. SYSTEM DESIGN

We first briefly introduce background information on coding and decoding schemes employed in this work, and then present our method to enable intra-frame bidirectional transmission.

We use intensity modulation for data transmission, which is also adopted by the IEEE 802.15.7 standard [2] developed for short-range communication using visible light of wide-band light bulbs. Binary information is mapped to the presence (symbol HIGH) or absence (symbol LOW) of the visible light carrier. The main PHY layer of the IEEE 802.15.7 standard uses the OOK modulation with the Manchester Run-Length Limited (RLL) line code. RLL line codes are used to avoid long runs of light on and light off that could end up in flicker effects, as well as clock and data recovery detection problems. Therefore, bit 1 is mapped to symbol sequence LOW-HIGH, and bit 0 is mapped to HIGH-LOW. Demodulation is performed with direct detection. Based on the received signal’s voltage, the receiving node detects the received signal as the sequence of symbols HIGH and LOW that are then converted to binary data.

2.1 Intra-Frame Bidirectional Transmission

The key enabler of intra-frame bidirectional transmission is that a node normally does not need to emit light when transmitting a symbol LOW. Therefore, the transmission of data is partially disentangled from the actual operation of the node, an essential difference from RF communication.

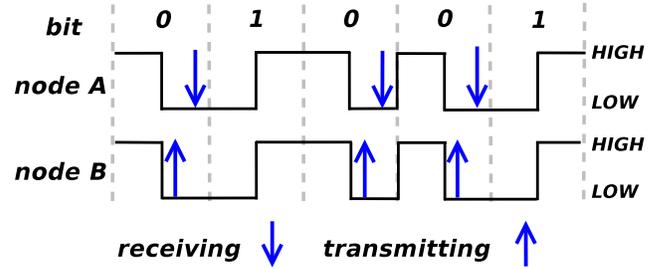


Figure 2: Key technique for intra-frame bidirectional transmission: node A can receive one symbol when transmitting the symbol LOW of bit “0” of the line code.

Let us consider two nodes, A and B. We assume that node A has gained access to the medium and it starts to transmit data to node B. *Since node A is “idle” when it transmits a symbol LOW, it can make use of this time to receive a data symbol.* To cooperate with this, node B can start itself to transmit a symbol if it can *predict* that it will receive a symbol LOW in the next symbol slot.

Under the assumption of Manchester RLL code, the prediction is based on: *i*) node B has received a HIGH symbol in current symbol slot; *ii*) the received HIGH symbol is the first part of a modulated bit (i.e., bit “0”). We illustrate an example in Fig. 2. Just with one LED, Node B is able to send data during the reception of a frame, while node A can receive data when it transmits a frame. This allows to enable intra-frame bidirectional transmission.

As shown in Fig. 2, node A should expect to receive a symbol only when it transmits the LOW symbol of bit “0”. The reason is that node B can only predict the LOW symbol of bit “0”, but it can not predict the LOW symbol of bit “1”. In expectation, half of data from node A are with bit “1” and the other half with bit “0”. Therefore, node B will transmit data for half of the payload of node A. This comes at the minor cost of not performing clock and data recovery for half of the data, which can be accepted for reasonable clock drift.

The approach can be extended to other RLL line codes. In more general terms, the intra-frame bidirectional transmission can be enabled for those symbols where there is a probability of one to jump to a state where node A transmits the LOW symbol.

2.2 The CSMA/CD-HA MAC Protocol

The Carrier Sensing Multiple Access / Collision Detection&Hidden Avoidance (CSMA/CD-HA) protocol is proposed to ensure fair channel access among all VLC nodes and reduce the impact of collisions and hidden nodes. We introduce our protocol, referring to the notations given in Table 1. When a frame is available for transmission, the MAC first senses the channel. The frame is transmitted immediately if the channel is sensed clear. If the channel is assessed busy, the MAC starts a backoff counter. The frame is transmitted when the counter reaches zero. The transmitting node and its corresponding receiver become a P-TX and a P-RX, respectively. Upon frame transmission, the P-TX switches to receiving mode when it transmits the LOW symbols of bit “0”. The P-RX, similarly, switches to transmit a symbol when it finishes to receive a HIGH symbol of bit “0” (thus predicting that the next incoming symbol is LOW), to

Table 1: Notations used in the CSMA/CD-HA MAC protocol.

Primary TX (P-TX)	Node that has gained access to the medium with backoff protocol
Primary RX (P-RX)	Peer RX of a P-TX
Secondary TX (S-TX)	Node (not a P-RX) transmitting an in-band frame with a P-TX’s transmission
Secondary RX (S-RX)	Peer RX of a S-TX
Primary frames	Frames sent by a P-TX to a P-RX
Embedded frames	Frames sent opportunistically by either a P-RX or a S-TX

perform the intra-frame bidirectional transmission, as presented in Sec. 2.1.

After the P-TX successfully receives the embedded frame, it appends an-octet ACK to the frame it is transmitting to acknowledge the reception of embedded data. The P-RX sends an ACK to the P-TX when it decodes a frame without error. The whole process is illustrated in Fig. 3 (with node A as the P-TX and B as the P-RX). If the P-TX has not received an ACK within the timeout, a retransmission occurs and the contention window is doubled unless it reaches the maximal contention window. The frame is dropped when a pre-defined number of retransmissions fail.

Hidden node avoidance As a result of intra-frame bidirectional transmission, the P-RX “always” sends signals when it is receiving a frame, and part of them are HIGH symbols. These HIGH symbols can be sensed by the nodes under its coverage, thus they will not send frames to the receiver during this period and the potential hidden-node problem is alleviated. The only part of the frame which is not protected to hidden nodes is before the embedded frames start to be transmitted, since during this time, a hidden node to the TX may judge the channel is clear. However, this duration is usually short compared to the whole frame transmission period and the influence tends to be very low. This is similar to the RTS/CTS handshake (which sends short frames to nodes under the TX and RX’s coverage, to announce that the shared channel will be occupied for some time.) Some-

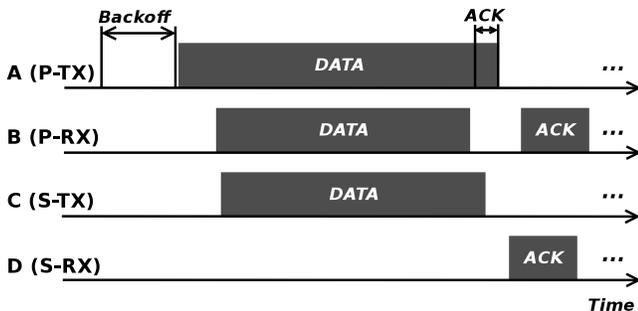


Figure 3: Proposed MAC protocol that supports intra-frame bidirectional transmission: an illustration for the left scenario of Fig. 1. (Note: the DATA/ACK frames shown here are transmitted by their corresponding nodes, not received.)

times the P-RX may have no data to transmit. Under this case, the P-RX can send dummy symbols to announce that the channel is busy while it is receiving a frame. To summarize, the embedded frames of the primary communication have twofold use: send additional in-band data and act as active acknowledgement of ongoing primary data reception.

Collision detection Traditional collisions caused by nodes reaching the backoff counter equal to zero are still possible. These collisions are immediately detected by the absence of embedded data during the payload, and thus the transmissions are stopped without waiting for the end of the payload.

Secondary intra-frame concurrent transmissions Further intra-frame concurrent transmissions can also be enabled, depending on the number of existing nodes, the FOVs of their LEDs and their positions. Take the left scenario in Fig. 1 as an example. The wide-FOV LEDs of node A make it possible to communicate with nodes B, C, and D, while the latter three have limited connectivity due to their narrow-FOV LEDs. We illustrate an example of the protocol in Fig. 3. Nodes C and D can receive A’s signals, and C can detect that it is not (as well as D) the intended receiver of data from A. Node C then transmits embedded frames to D as node B transmits to A (C thus becomes a S-TX), while D can receive these frames following exactly A’s receiving steps (D thus becomes a S-RX). After successfully decoding a frame from node C, D sends back an ACK without causing any interference to the communication between nodes A (the P-TX) and B (the P-RX).

2.3 Discussion

We continue to discuss some other important aspects of our MAC protocol.

Length of an embedded frame The length of an embedded frame heavily depends on the number of bit “0” containing in the primary frame body. This number can be calculated by the transmitter and appended to the MAC frame header.

When to send an embedded frame An embedded frame should be transmitted after a node receives the source address of the current primary frame. Only the intended receiver of the transmission can be declared as P-RX. Other nodes may act as S-TX, as long as they are not the intended receivers of P-RX. They start a secondary backoff process that follows the same rules as the primary one to decide the access to the medium.

Robustness of intra-frame concurrent transmissions

The intra-frame current transmissions should be robust to expected or unexpected events, i.e., the secondary transmission should be disabled immediately if the primary transmission stops due to detections of collisions or hidden nodes or other unexpected errors. The CSMA/CD-HA protocol has this ability. For example, an underlying assumption of secondary transmissions is that the FOVs of nodes is such that the transmission of S-TX and S-RX does not interfere with the transmission of P-TX and P-RX. If this is not the case, the primary transmission is stopped due to detected collisions. As soon as P-TX stops its transmission because of a detected collision, the secondary transmission of S-TX and S-RX will be disabled for a while.

3. IMPLEMENTATION

We implement the proposed MAC protocol in our general-purpose software-defined open source platform OpenVLC for

visible light communication networks [14]. The MAC protocol is implemented as a Linux driver – another network interface – that is able to communicate directly with the hardware and the Linux networking stack.

3.1 Hardware

The platform consists of a BeagleBone Black (BBB) board [1] and a front-end transceiver. The BBB board is equipped with the TI AM3359 CPU (1GHz), 512MB DDR3 RAM and 2GB on-board flash storage, and provides 4 timers and 65 General-Purpose Input/Output (GPIO) pins [1]. The front-end transceiver employs a single LED to both transmit and receive along with a few basic electronic components. This design has been proved to be resilient to ambient light interference without additional electronic processing [6] and we envision that it can be well suited to enable a communication network of consumer devices with LED front-end.

We rely on basic off-the-shelf electronic components and only use a very basic Physical Layer (PHY), though it can be scaled to use more advanced PHYs. The block diagram of the transceiver is shown in Fig. 4. It includes a TransConductance Amplifier (TCA) for transmission, a TransImpedance Amplifier (TIA) and an Analog-to-Digital Converter (ADC) for reception, a tristate-output buffer, and ancillary circuitry for transmission and reception.

A prerequisite to enable the intra-frame bidirectional transmission method is that nodes can switch between transmitting mode and receiving mode on a symbol basis. In our implementation, a software-defined TX/RX controller is used to switch the LED between TX and RX through the GPIO pins: *i*) in **TX mode**, encoded signals are amplified by the TCA and then fed to the forward-biased LED; *ii*) in **RX mode**, light signal is received by the reverse-biased LED and then the received small photocurrent is amplified by the TIA. Finally, an ADC converts the output analog signals to digital signals, which are then sent to the decoder through the Serial Peripheral Interface (SPI).

Through the TX/RX controller and the tri-state buffer, the LED can be switched between being TX mode and RX mode with low latency. This makes it possible to enable intra-frame bidirectional transmission as presented in Sec. 2.1.

3.2 Software

The Linux operating system running within the BBB board is the Angstrom Distribution with kernel version 3.8.13. Fig. 5 illustrates the software stack of the platform, where the MAC is built based on primitive functions we implement. These

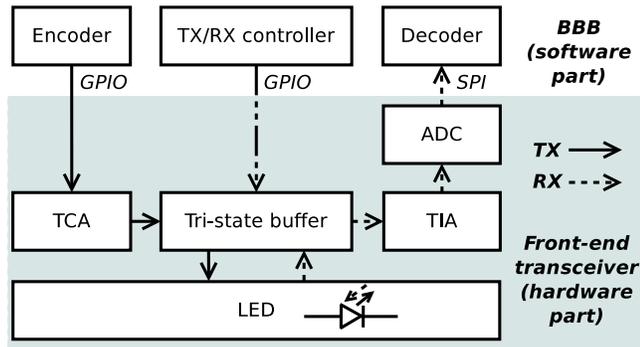


Figure 4: Diagram of the VLC platform.

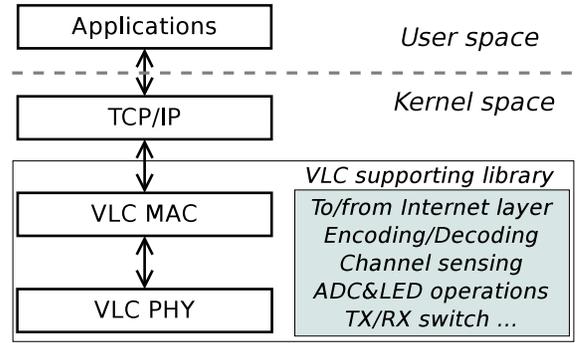


Figure 5: Implementation of the MAC protocol in embedded Linux operation system.

functions include writing a symbol to the LED, reading a symbol from ADC, coding/decoding, preamble detection, and so on. Furthermore, we implement functions to receive and send packets to the Internet layer. In this way, the MAC is connected with the TCP/IP layers of the Linux operating system and thus the MAC becomes transparent to various applications.

In our implementation, there are two types of frames in the MAC layer, DATA and Acknowledgement (ACK). The frame structure and octets each field occupies are shown in Fig. 6. Distinguishing DATA and ACK frames is through the length of frame body (payload): if the frame has no payload (i.e., “Length”= 0), it is an ACK. Otherwise, it is a DATA frame. Each DATA frame can carry a payload up to 255 bytes. The MAC destination and source addresses that follow the “Length” field each occupies 2 bytes. The 2-byte field “Protocol” identifies the upper layer protocol encapsulated in the payload of current frame. The fields from “Length” to (upper layer) “Protocol” form the MAC frame header. A two-octet Cyclic Redundancy Check (CRC) over the frame header and frame body is appended to the end of each frame. A three-octet preamble is added to the beginning of each frame for frame synchronization.

In order to enable embedded communication, for simplicity we currently assume that the number of bit “0” is estimated by the receiver based on the length of primary frame body and the assumption that bits “0” and “1” are evenly distributed.

4. EVALUATION

In this section we evaluate the proposed CSMA/CD-HA protocol through experiments. The symbol slots of symbols LOW and HIGH are both set to 200 μ s. An 8-symbol period is used for channel sensing, i.e., to determine whether the channel is clear or not. The minimum and maximum value of the contention window are set to 4 and 32, respectively. The maximal frame retransmission time is set to 4. Different to a primary frame, an embedded frame is coded with OOK modulation without the Manchester code (bit “0” is mapped

Preamble	SFD	Length	Dst	Src	Protocol	Payload	CRC
3 Bytes	1 B	1 B	2 B	2 B	2 B	0-255 B	2 B

Figure 6: Frame format of DATA and ACK: Length > 0 \iff DATA; Length= 0 \iff ACK.

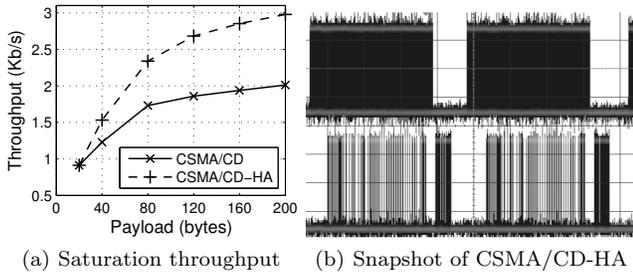


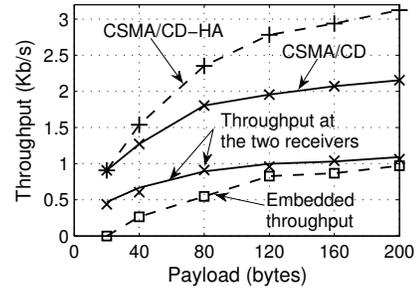
Figure 7: Evaluation results of a point-to-point link.

to symbol LOW and bit “1” to symbol HIGH.). Details of the electronic devices adopted in the testbeds can be found in [14]. Note that we use two narrow-FOV LEDs to emulate a wider-FOV LED when needed.

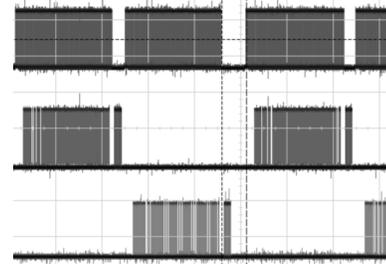
We compare CSMA/CD-HA with the CSMA/CD introduced in [6]. CSMA/CD uses the same backoff mechanism of the traditional Carrier Sensing Multiple Access/Collision Avoidance (CSMA/CA). While transmitting, a node adopting CSMA/CD senses the channel when it transmits a symbol LOW and stops the current transmission if the channel is declared busy. It has been shown in [6] that the CSMA/CD outperforms the CSMA/CA.

Point-to-point link We evaluate the MAC layer saturation throughput of a point-to-point link under the CSMA/CD-HA and CSMA/CD protocols, where two nodes are within the FOV of each other. The saturation throughput is achieved under the setting that both nodes always have data to transmit to each other. Throughput versus per-frame payload is shown in Fig. 7 (a), where the payload varies from 20 to 200 bytes. As expected, the throughputs under both protocols increase as the payload increases, e.g., under CSMA/CD, the saturation throughput is 1 kb/s when the payload is 20 bytes, while the throughput achieves a value of 2 kb/s when the payload increases to 200 bytes. Another observation from Fig. 7 (a) is that our proposed CSMA/CD-HA greatly outperforms CSMA/CD when the payload is long, e.g., when the payload is 200 bytes, the throughput under CSMA/CD-HA is about 150% of that under the CSMA/CD. This achievement comes from intra-frame bidirectional transmission in CSMA/CD-HA. The throughputs under the two MAC protocols are the same when the payload is 20 bytes. This is because the transmission of an embedded frame becomes impossible when the payload of a primary frame is too short. An oscilloscope snapshot of the protocol is shown in Fig. 7 (b), where we can see clearly the interaction of nodes and the transmission of embedded frames.

Three-node network without hidden nodes The performance of a three-node network without hidden nodes is evaluated, as shown in Fig. 8 (a). In this scenario, one node acts as a base station and keeps sending data alternately to the other two nodes (i.e., the base station uses the round-robin scheduling algorithm). For instance, this scenario may emulate a setup where the base station is the light bulb in the ceiling and most of the traffic is downlink. As expected, we observe from the figure that the throughputs at the two receivers are similar under various frame payloads. Another observation is that the throughput gain of CSMA/CD-HA over CSMA/CD increases with the frame payload. This gain achieves a value of 50% when the frame payload is 200 bytes.



(a) Saturation throughput



(b) Snapshot of CSMA/CD-HA

Figure 8: Evaluation results of the three-node network without hidden nodes.

The interactions between these three nodes are illustrated by the oscilloscope snapshot in Fig. 8 (b).

Three-node network with hidden nodes We also evaluate a three-node network with hidden nodes, similar to the illustration shown in the right sub-figure of Fig. 1. We compare the CSMA/CD-HA with the CSMA/CD without RTS/CTS. The position where a collision is detected within each collided frame is shown in Fig. 9. We can observe from Fig. 9 (a) that under CSMA/CD, most of the collisions happen at the beginning of frames. This is because the two TXs can not hear each other’s transmission, thus causing a collision when both of them start transmitting. However, the collisions under CSMA/CD-HA occur either at the beginning or at the end of collided frames, as shown in Fig. 9 (b). This is because a RX starts sending an embedded frame after it receives the source address of the corresponding primary frame, and finishes sending the embedded frame before the whole primary frame is fully transmitted (as illustrated in Fig. 3). Although there are still collisions caused by the hidden node under CSMA/CD-HA, the number of collisions is actually reduced greatly. This is shown from Fig. 9 (c), where the saturation throughput versus frame payload is presented. We can observe that CSMA/CD-HA always outperforms CSMA/CD greatly independent of frame payloads, and the throughput improvement is up to 1000% when frame payload is set to 200 bytes.

Four-node network The performance of a four-node network is also evaluated. The setting of this scenario can be referred to the left sub-figure of Fig. 1. The four nodes act as a P-TX, P-RX, S-TX and S-RX, respectively. The oscilloscope snapshot of the interactions between these nodes is presented in Fig. 10 (a). We can see that three intra-frame concurrent transmissions are obtained. The corresponding saturation throughputs are presented in Fig. 10 (b). We observe that with the increase of frame payload, the throughputs under

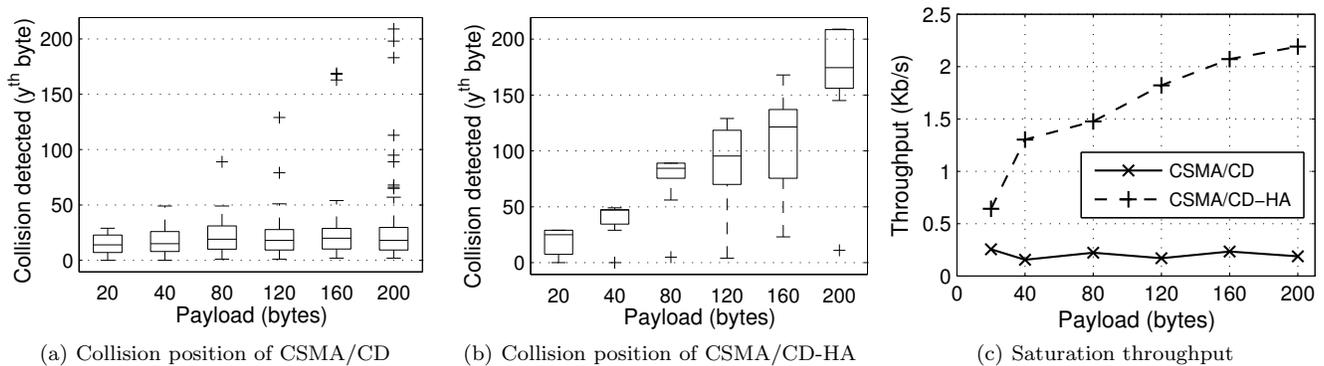


Figure 9: Evaluation results of the three-node network with hidden nodes.

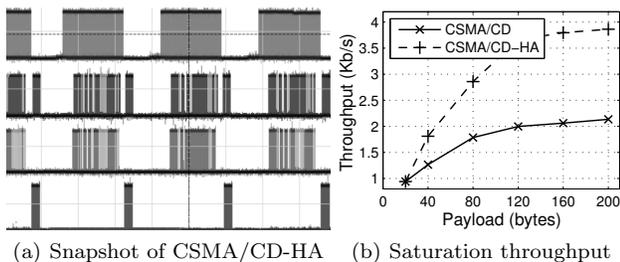


Figure 10: Evaluation results of a four-node network.

both protocols increase. The CSMA/CD-HA starts to outperform CSMA/CD after the payload is more than 20 bytes. The throughput gain under CSMA/CD-HA over CSMA/CD increases to around 100% when the payload is 200 bytes. This gain is twice as large as those under the scenarios of point-to-point link and three-node network without hidden nodes. The reason for this is in the four-node network, an additional intra-frame transmission is enabled between the S-TX and S-RX.

5. RELATED WORK

VLC has received strong attention from the designers of the next generation of cellular networks [13] and the point-to-point communication using smartphones [7, 8] and cars [11, 15]. The IEEE has also developed the 802.15.7 standard [2] for short-range communication using visible light.

Single antenna VLC communication While photodiodes are normally used as receivers, a reverse-biased LED instead of a photodiode is used in [5] as a receiver to implement a bidirectional communication network. This principle has been exploited by [6] to introduce a LED-to-LED communication network and is further developed in [12]. All these works operate on microcontrollers and are implemented as embedded software in the environment of non-operating system. Connecting these works with various networking protocols is not straightforward. In contrast, our implementation resides in the Linux operating system and appears as a normal network interface that can be easily extended to support different upper layer protocols.

Full-duplex RF/VLC communication In-band full-duplex RF communication was proposed in [4] and further implemented as a prototype in [9]. The key technique here is a device that uses the inverse of its transmitted signals to

cancel the self-interference to its received signal. Compared to our proposed technique, both of them can achieve in-band bidirectional transmission. A different is that [4, 9] used two antennas for transmitting and receiving, respectively, while in our technique, the same LED is used for both transmitting and receiving purposes. Recently, a method combining analog and digital cancellation is proposed in [3]. This method enables the full duplex communication with a single antenna. To the best of our knowledge, there is yet no work on full-duplex (in-band) VLC, although some authors have proposed the use of visible light and infrared light or LEDs operating on different wavelengths for “full-duplex” communications, as in [10].

6. CONCLUSION

In this paper, we presented the design, implementation, and performance evaluation of the CSMA/CD-HA protocol for communication networks using visible light. By exploiting the OOK modulation used in VLC systems, we enabled intra-frame bidirectional transmission for networks of LEDs with heterogeneous FOVs. We implemented CSMA/CD-HA in an embedded platform and showed its superior performance for the detections of collision and hidden transmissions with respect to recent works in visible light communication networks. We expect that the next challenge to implement CSMA/CD-HA in real systems will be to study how well the protocol works at faster switching rate between symbols HIGH and LOW. For instance, this requires the design of novel synchronization schemes at system level, necessary for the transmission of data while receiving, and the understanding of the impact of the clock drift caused by not using half of the symbols of the line code for clock and data recovery.

7. REFERENCES

- [1] BeagleBone Black. <http://beagleboard.org/Products/BeagleBone+Black>.
- [2] IEEE standard for local and metropolitan area networks—part 15.7: Short-range wireless optical communication using visible light. *IEEE Std 802.15.7-2011* (Sept 2011), 1–309.
- [3] BHARADIA, D., McMILIN, E., AND KATTI, S. Full duplex radios. In *Proceedings of the ACM SIGCOMM* (2013), pp. 375–386.
- [4] CHOI, J. I., JAIN, M., SRINIVASAN, K., LEVIS, P., AND KATTI, S. Achieving single channel, full duplex

- wireless communication. In *Proceedings of the ACM MobiCom* (2010), pp. 1–12.
- [5] DIETZ, P., YERAZUNIS, W., AND LEIGH, D. Very low-cost sensing and communication using bidirectional LEDs. In *TR2003-35* (2003).
- [6] GIUSTINIANO, D., TIPPENHAUER, N., AND MANGOLD, S. Low-complexity visible light networking with LED-to-LED communication. In *Proceedings of the IFIP Wireless Days (WD)* (2012), pp. 1–8.
- [7] HAO, T., ZHOU, R., AND XING, G. Cobra: Color barcode streaming for smartphone systems. In *Proceedings of the ACM MobiSys* (2012), pp. 85–98.
- [8] HU, W., GU, H., AND PU, Q. Lightsync: Unsynchronized visual communication over screen-camera links. In *Proceedings of the ACM MobiCom* (2013), pp. 15–26.
- [9] JAIN, M., CHOI, J. I., KIM, T., BHARADIA, D., SETH, S., SRINIVASAN, K., LEVIS, P., KATTI, S., AND SINHA, P. Practical, real-time, full duplex wireless. In *Proceedings of the ACM MobiCom* (2011), pp. 301–312.
- [10] LIN, K., AND HIROHASHI, K. High-speed full-duplex multiaccess system for LEDs based wireless communications using visible light. In *Proceedings of the International Symposium on Optical Engineering and Photonic Technology (OEPT)* (2009), pp. 1–5.
- [11] LIU, C. B., SADEGHI, B., AND KNIGHTLY, E. W. Enabling vehicular visible light communication (V2LC) networks. In *Proceedings of the ACM VANET* (2011), pp. 41–50.
- [12] SCHMID, S., CORBELLINI, G., MANGOLD, S., AND GROSS, T. R. LED-to-LED visible light communication networks. In *Proceedings of the ACM MobiHoc* (2013), pp. 1–10.
- [13] TSONEV, D., VIDEV, S., AND HAAS, H. Light Fidelity (Li-Fi): towards all-optical networking. vol. 9007, pp. 1–10.
- [14] WANG, Q., GIUSTINIANO, D., AND PUCCINELLI, D. OpenVLC: Software-defined visible light embedded networks. In *1st ACM Workshop on Visible Light Communication Systems , in conjunction with MobiCom 2014* (September 2014), pp. 1 –6.
- [15] YU, S.-H., SHIH, O., TSAI, H.-M., AND ROBERTS, R. Smart automotive lighting for vehicle safety. *IEEE Communications Magazine* 51, 12 (2013), 50–59.