LightTour: Enabling Museum Audio Tour with Visible Light

Lennert Vanmunster
KU Leuven, Belgium
lennertvanmunster@gmail.com

Jona Beysens
KU Leuven, Belgium
jona.beysens@kuleuven.be

Qing Wang
TU Delft, the Netherlands
qing.wang@tudelft.nl

Sofie Pollin
KU Leuven, Belgium
sofie.pollin@kuleuven.be

Abstract

We exploit visible light in museums to design LightTour, a wirelessly networked system empowered by visible light that provides audio tour service to visitors. LightTour supports full-duplex transmissions with a single multi-band optical receiving antenna at each transceiver: for the downlink, the access points (APs) broadcast audio streams to visitors' devices through visible light communication; for the uplink, the devices send visitor's feedback to APs through infrared light transmission. To network a number of APs and devices, we identify two main challenges: (1) self-interference at the APs, caused by the floor reflection of their downlink signals, that degrades uplink transmissions; (2) the uplink hidden-device problem. In this work, we develop an online low-complexity method that can achieve maximally 30 dB self-interference cancellation at APs. We design a new MAC protocol – full-duplex carrier sense multiple access with collision detection & hidden avoidance (FD-CSMA/CD-HA) – to solve the hidden-device problem and to provide reliable and high-throughput audio broadcasting. To evaluate the system performance of LightTour, we build a prototype with off-the-shelf components. Besides, we develop a simulator to assess large-scale scenarios. The results validate our self-interference cancellation method's feasibility and demonstrate FD-CSMA/CD-HA's robustness and its advantage over existing protocols.

1 Introduction

Audio tours have become ubiquitous in museums. An audio tour provides a visitor with location-based spoken commentary about museums' exhibition through a hand-held device. In this way, visitors can experience museums in a more personalized and effortless manner. Modern audio tour systems can be broadly divided into four categories based on their operating principles:

1. *Local audio storage with manual track selection:* audio is stored locally on each device. These systems require user input, usually via buttons or a touchscreen, for audio track selection and playback.
2. *Local audio storage with location-aware track selection:* these systems sense the user's location using radio-frequency (RF) based positioning methods. The audio track stored on the hand-held device is played back based on the current user position.
3. *Wireless audio transmission with manual track selection:* in these systems the user manually selects the audio track on the device. Subsequently, the device sends a request for the track to a nearby AP, which then transmits that audio track to the device using RF communication.
4. *Wireless audio transmission with location-aware track selection:* these systems are the most advanced and include both a mechanism for user localization and wireless audio transmission through RF communication.

Motivated by the pervasive existence of visible light in museums and the recent advances in visible light communication (VLC), we design an indoor location-based audio tour system empowered by light. In our system *LightTour*, automated location-based audio transmission is realized by modulating the LEDs whose primary function is to provide illumination (Category 4 system above). Additionally, LightTour supports *full-duplex transmissions* in which an uplink channel is provided for users to send real-time feedback messages to the museum. The uplink channel for real-time, location-based feedback messages can improve future user experience and museum management. Another example of an uplink application that improves museum management is real-time tracking and collection of visitor locations at a central entity.

The advantages of light communication enabled audio tour systems over existing RF-based systems are threefold. The first advantage is the reduced installation cost and hardware complexity. Not only do RF systems require additional installation of APs, they also increase the transmitter/receiver (TX/RX) hardware complexity significantly. LightTour, on the other hand, offers easy installation and low hardware complexity by integrating simple TX/RX hardware into the lighting infrastructure and exploiting a single

International Conference on Embedded Wireless Systems and Networks (EWSN) 2021
17–19 February, Delft, The Netherlands
ISBN: 978-0-9949886-5-2

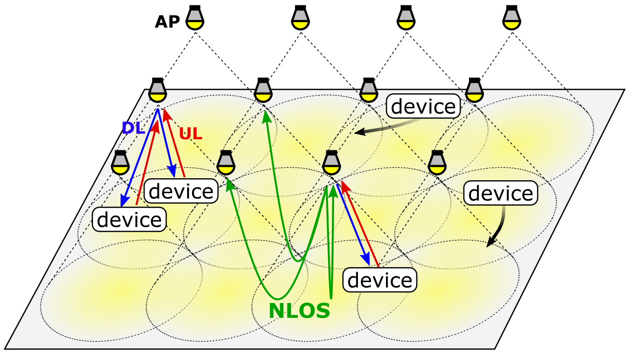

Figure 1. Network topology (UL: line-of-sight uplink; DL: line-of-sight downlink)

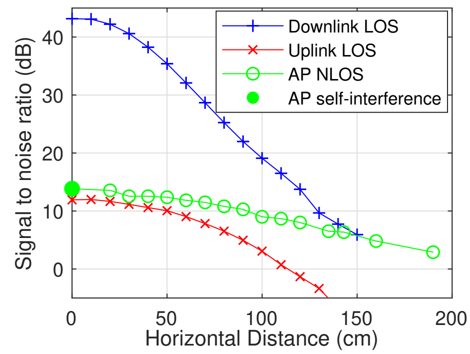

Figure 2. Measured strength of different signals (cf. Fig. 9 for the experimental setup)

antenna for multi-band reception. Secondly, by combining illumination and communication, LightTour inherits the advantage of VLC systems, being very energy-efficient [20]. Thirdly, user localization methods in RF systems, e.g. based on received signal strength indication (RSSI), time of arrival, angle of arrival or fingerprinting, suffer from multipath fading, degrading their accuracy and reliability [19]. Additionally, those methods require more APs than would be needed for data service alone, increasing interference, overall cost and complexity. Furthermore, the large-scale (large building with many rooms) and dynamic (movement of people) museum environment further exacerbates the multipath problem. In addition, RF localization methods suffer from other problems limiting their practicality, such as stringent AP time synchronization (time of arrival), high hardware complexity (angle of arrival) or significant system calibration (fingerprinting). With VLC, on the other hand, the inherent small cell size originating from its high directionality and blocking by objects, facilitates a simple and robust line-of-sight (LOS) proximity-based service.

Because of the above advantages of using light to design museum tour systems, already several LightTour-like but simpler systems are designed and deployed in reality, for example, the tour system with visible light in Pompeii of Italy and the tour system with infrared light in the Peace Palace of the International Court of Justice, Netherlands.

In this work, we target to design a practical and full-duplex wireless communication system with light for a network of APs and devices, and to solve potential challenges. To support full-duplex communication, VLC systems usually use visible light for the downlink and infrared light for the uplink [3]. Most of them only consider single-cell communication, i.e., an AP and one or several devices. Wireless networking between several APs and a number of devices through light, as shown in Fig. 1, has rarely been considered. Besides, going beyond the state of the art, we target at a system where *only a single multi-band optical sensor (i.e., receiving antenna)*, capable of sensing both visible and infrared light, is equipped at each AP and each device. This reduces hardware cost as no (expensive) optical filter and only one photodiode, filter circuit and ADC is required at each AP and device. To achieve such a practical system, two main challenges must be addressed:

Challenge 1: Non-line-of-sight (NLOS) signal. We observe from our experimental measurements in Fig. 2 that the strength of the AP self-interference, caused by floor reflections of its own downlink signals, is similar to the signal strength of uplink signals. Therefore, the NLOS self-interference signal degrades full-duplex uplink transmissions. However, the inter-AP NLOS signal also offers the opportunity for neighbouring APs to sense each other's transmissions (cf. green curve in Fig. 2) and allows the medium access control (MAC) protocol to achieve an efficient downlink channel utilization[1] because it avoids the hidden-AP problem. Thus, the challenge is how to exploit the reflected NLOS signal to improve the network performance.

Challenge 2: Hidden-device problem. To decrease the energy consumption at the devices and for eye safety reasons [25], the output power of the infrared LEDs should be much lower than that of the white LEDs at the APs[2]. To reach a sufficient infrared LED radiant intensity for uplink transmission, the half power semi-angle of infrared LEDs should be limited[3]. The low output power and small half power semi-angle mean that the devices are almost never able to sense each other's uplink transmissions, creating the hidden-device problem. Furthermore, when the devices are held by visitors, they can easily block the LOS link between neighboring devices, making the hidden-device problem more pronounced.

Contributions. In this work, we design, implement and evaluate the LightTour system which serves audio data to a large number of users based on full-duplex links with visible light and infrared light. Below we list our main contributions:

Contribution 1: A low-complexity digital self-interference cancellation (SIC) mechanism. We develop a VLC SIC mechanism that achieves up to 30 dB of cancellation in real-time on a microcontroller. SIC is realized in a very low-complex way using a lookup-table combined with real-time

[1] Channel utilization is defined as the fraction of time the channel is used for successful transmissions [36].

[2] The transmission power of white LEDs at the APs should be high enough to provide the required indoor illumination level.

[3] With a larger half power semi-angle, more transmission power is required in order to reach the same vertical communication distance.

path loss estimations. (Sec. 3)

Contribution 2: A full-duplex carrier sense multiple access with collision detection and hidden avoidance (FD-CSMA/CD-HA) MAC protocol. Based on the proposed SIC mechanism, which enables full-duplex transmissions, we design the MAC protocol FD-CSMA/CD-HA. Thanks to collision detection (CD), FD-CSMA/CD-HA provides reliable, high-throughput broadcasting of audio frames, without requiring any additional wired/wireless communication techniques. CD is enabled by full-duplex transmissions with light, made possible by our SIC mechanism. Furthermore, inspired by busy tone multiple access (BTMA) in RF [5], the hidden-device problem in LightTour is solved by periodically inserting ultra short 'busy slots' during downlink data frame transmission to indicate the uplink channel status. The 'busy slot' mechanism can be interpreted as an adaptation of BTMA to single carrier modulation schemes commonly employed in VLC. (Sec. 3)

Contribution 3: Proof-of-concept implementation. A proof-of-concept prototype is implemented with simple transmitter and receiver front-ends and the low-cost Arduino Due micro-controllers. (Sec. 4)

Contribution 4: Performance and robustness evaluation. Our results validate the feasibility of the SIC mechanism in practice and show the advantages of FD-CSMA/CD-HA over existing protocols. For example, the downlink channel utilization at high downlink load is increased by 185% compared to the IEEE 802.15.7 CSMA/CA protocol [1]; the uplink data rate is improved by 146% in high uplink load scenarios. Further, we evaluate the system's robustness. Our results demonstrate that blockage and ambient light can negatively impact the performance of our protocol. However, these negative impacts can be largely reduced or even eliminated by properly selecting the parameters related to the clear channel assessment (CCA) and CD mechanisms. (Sec. 5 and Sec. 6)

2 System Overview

The system architecture of LightTour is shown in Fig. 1. It consists of a number of APs, each with an LED transmitter and a single photodiode receiver. The APs are mounted on the ceiling and provide both illumination and communication. The MAC protocol does not require centralized synchronization among the APs, which reduces system complexity and improves scalability.

Similar to a traditional audio tour system, in LightTour each visitor is equipped with a hand-held device, through which received audio is played back. The device contains a photodiode receiver and infrared LED transmitter for communication with APs. Downlink and uplink communication is realized via a LOS link between the devices and their closest APs. Only a single photodiode, sensitive to both infrared light and visible light, is used at each AP and device.

Each AP stores one audio track and broadcasts audio streams to the devices in its range, enabling a location-based service. Broadcasting allows each AP to serve a large number of devices with a low physical layer (PHY) data rate.

On the uplink, the device sends visitors' feedback on the museum exhibition to the APs. The feedback can either be

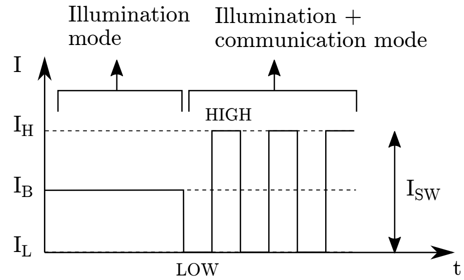

Figure 3. LED operating mode

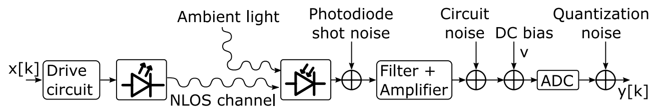

Figure 4. The model of the reflected NLOS visible light communication link

short text messages or audio pieces. More details about the application layer are presented in Sec. 3.3.

PHY modulation and coding. The used PHY modulation is based on the IEEE 802.15.7 standard [1]. The downlink uses on-off-keying (OOK) with 4b6b line coding. In addition, a Reed-Solomon code (168,200) is used to improve the transmission reliability. The uplink uses OOK with Manchester coding.

Each LED has two operating modes: 1) illumination mode, and 2) illumination+communication mode, as shown in Fig. 3. When an LED is in illumination mode, a constant bias current I_B flows through the LED to achieve the desired illumination. When it operates in illumination+communication mode, the light intensity emitted from the LED is modulated to transmit data using OOK with 4b6b line coding, in which a swing level I_{sw} around the bias is adopted to represent the LOW and HIGH symbols. A current $I_L = 0$ represents a LOW symbol, while a current $I_H = I_{sw}$ denotes a HIGH symbol.

3 Protocol Design

We first introduce the low-complexity SIC mechanism that enables full-duplex transmissions. Then, we present the MAC protocol. Finally, we describe the application layer.

3.1 Self-Interference Cancellation (SIC)

To enable full-duplex transmissions with a single receiving antenna at APs, we develop a low-complexity SIC mechanism. It operates in real-time on the microcontroller at each AP.

3.1.1 Problem with existing SIC algorithms

The NLOS link model is shown in Fig. 4. The received signal sample $y[k]$ at time instance k, *which is the reflected self-interference signal*, can be written as (bold symbols indicate vectors) [14]:

$$y[k] = \boldsymbol{h}[k] * \boldsymbol{x}[k] + d + n[k], \quad (1)$$

with "$*$" the convolution operator, $\boldsymbol{h}[k]$ the overall system impulse response (including channel impulse response and circuit impulse response), $\boldsymbol{x}[k]$ a vector containing the last L transmitted symbols $x[k-L+1]$ to $x[k]$, in which $x[k]$ can be either -1 (LOW symbol in Fig. 3) or 1 (HIGH symbol) and L the length of the impulse response $\boldsymbol{h}[k]$. Next, $n[k]$ is the overall system noise (including photodiode shot noise, analog-to-digital converter (ADC) quantization noise, circuit thermal noise) and d the static direct current (DC) bias added in the receiver circuit. The DC bias is added such that the voltage at the ADC input falls approximately in the center of the ADC's input voltage range.

In state-of-the-art digital SIC mechanisms, used in RF communication, $\boldsymbol{h}[k]$ is usually estimated adaptively on a training sequence by using a recursive algorithm such as the least-mean-squares algorithm [22]. However, from our measurements in Table 2 of Sec. 5, we observe that in embedded networks the computation time of least-mean-squares (update $\boldsymbol{h}[k]$ and compute output $y[k]$) and filter evaluation (compute $y[k]$) are 3.40 μs and 1.56 μs per symbol respectively. The targeted symbol duration of our system is 2 μs, hence filter evaluation and least-mean-square calculation would severely limit achievable data rate.

3.1.2 *Our proposed low-complexity SIC*

We leverage the unique VLC channel characteristics to simplify the computations and achieve real-time SIC. We split the system impulse response $\boldsymbol{h}[k]$ in (1) into two parts: $\boldsymbol{h}[k] = \boldsymbol{h_r} \cdot g[k]$. The first part, $\boldsymbol{h_r}$, is a static vector (does not change over time) and is the normalized (DC gain=1) impulse response of the receiver circuit consisting of the filter, amplifier and ADC. The second part, $g[k]$, is a dynamic scalar (changes over time) and is the overall DC gain of the system which comprises the dynamic channel path loss $h_c[k]$, static photodiode sensitivity and static receiver amplifier gain. In reality, the VLC channel has a multiple-tap impulse response $\boldsymbol{h_c}[k]$. We approximate $\boldsymbol{h_c}[k]$ by a one-tap path loss $h_c[k]$, based on the observation that the root-mean-square delay spread of an indoor optical channel is ~ 20 ns [14] and thus $\boldsymbol{h_c}[k]$ is much shorter than the target symbol time in our system (2 μs). Then (1) becomes:

$$y[k] = \boldsymbol{h_r} * (g[k] \cdot x[k]) + d + n[k]. \tag{2}$$

To compute the SIC in real-time in a microcontroller, we replace $\boldsymbol{h_r}$ with a static lookup-table $\mathcal{T}$, trading time complexity for memory complexity. Let $x_u[k]$ be a unipolar representation (LOW=0, HIGH=1 instead of LOW=-1, HIGH=1) of $x[k]$ and let $\boldsymbol{x_u}[k]$ denote the binary concatenation of the vector that contains $x_u[k-L+1]$ to $x_u[k]$ (e.g., $\boldsymbol{x_u}[k] = 0b01010110011$ for $L = 11$). We refer to L as the lookup-table order which is similar to the number of filter taps of $\boldsymbol{h_r}$ replaced by $\mathcal{T}$. The AP then estimates $y[k]$ using the following formula:

$$\hat{y}[k] = \mathcal{T}_{\boldsymbol{x_u[k]}} \cdot \hat{g}[k] \cdot x[k] + \hat{d}, \tag{3}$$

in which $\hat{y}[k]$, $\hat{g}[k]$, and $\hat{d}$ denote estimates of $y[k]$, $g[k]$, and d, respectively. In addition, $\mathcal{T}_{\boldsymbol{x}_u[k]}$ denotes the element at index $\boldsymbol{x_u}[k]$ in the lookup-table. At each time k, the index $\boldsymbol{x_u}[k]$ is updated recursively in software by right shifting $x_u[k]$ into $\boldsymbol{x_u}[k-1]$:

$$\boldsymbol{x}_u[k] = (\boldsymbol{x}_u[k-1] \ll x_u[k]) \;\; \texttt{\&\&} \;\; 0x7FF \;\; (L = 11), \tag{4}$$

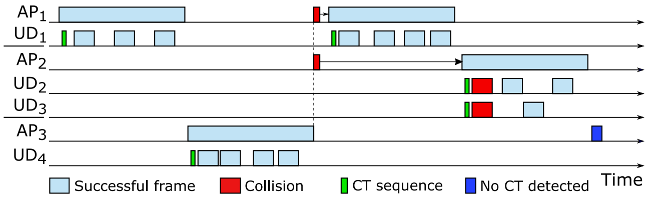

Figure 5. MAC timeline (CT: continue-transmission)

in which "$\ll$", "&&" and "0x" denote the left shift operator, logical 'and' operator and hexadecimal notation, respectively.

Estimate g and d. The least-square estimates $\hat{g}$ and $\hat{d}$ in (3) are computed in real-time on a random Manchester-encoded training sequence, transmitted at the beginning of each downlink frame (cf. Fig. 7), as follows:

$$\hat{g} = (\bar{y}^H - \bar{y}^L)/2; \quad \hat{d} = \bar{y} = (\bar{y}^H + \bar{y}^L)/2, \tag{5}$$

where $\bar{y}$ is the mean of $y[k]$ in the training signal, $\bar{y}^H$ and $\bar{y}^L$ are the mean of $y[k]$ during the transmissions of HIGH and LOW symbols, respectively.

Estimate h_r and $\mathcal{T}$. The static impulse response h_r and the lookup-table $\mathcal{T}$ are learned offline for each AP of the testbed in Fig. 9 using a least-squares method when transmitting random training data. Further measurements also confirm that $\mathcal{T}$ is sufficiently independent of the channel pass loss h_c.

The residual self-interference after cancellation $\rho[k]$ is calculated at each sample k as follows:

$$\rho[k] = y[k] - \hat{y}[k]. \tag{6}$$

To compute $\rho[k]$ per sample k from (3), (4) and (6) only a few logical operations (index update), additions and one multiplication are necessary. This leads to a decrease in computational complexity from $O(L)$ to $O(1)$. To store the lookup-table, 2^L bytes of memory are required, which is manageable for low values of L. We will analyze the performance and computation time of our low-complexity SIC with experiments in Sec. 5.

3.2 The FD-CSMA/CD-HA MAC protocol

Based on our SIC mechanism, we design a new MAC protocol: *full-duplex carrier sense multiple access with CD and hidden avoidance (FD-CSMA/CD-HA)*. We use Fig. 5 and Fig. 6 as illustration. Fig. 5 shows the timeline of the MAC protocol with 3 APs and 4 devices. The device UD_1 is located nearby AP_1, UD_2 and UD_3 are located nearby AP_2, and UD_4 is located nearby AP_3.

Downlink. Downlink transmissions are broadcast. The downlink part of FD-CSMA/CD-HA is based on CSMA/CD, where the main differences are the CCA and CD mechanisms adapted to VLC, and the CT sequence allowing APs to detect presence of nearby devices. High quality audio transmission requires near-zero frame loss. A broadcast solution that relies on ACK frames or request-to-send/clear-to-send (RTS/CTS) frames would be impractical. With broadcast transmissions and ACK and RTS/CTS, one device experiencing high interference drastically decreases overall system throughput since retransmissions are required for that device only. In addition, integrating ACK and RTS/CTS in a broadcast protocol is impractical since ACK and CTS frames transmitted from different devices would collide. The solution for broadcasting without frame loss due to collisions is CD at each AP. Our FD-CSMA/CD-HA provides reliable audio broadcasting since APs can detect detect and retransmit colliding frames. IEEE 802.15.7 CSMA/CA [1], on the other hand, has no CD, resulting in loss of colliding broadcast frames.

Clear channel assessment. Before transmitting any data, each AP assesses if the downlink channel is busy by measuring the energy of N_{cca} samples. If the energy is greater than a predefined threshold (due to the detected NLOS signal), then a busy channel is declared.

Collision detection. It is performed on a random, Manchester-encoded training and test sequence with length L_{ccd} bytes transmitted at the start of each frame. The training sequence is used for DC channel gain $\hat{g}$ estimation. During the transmission of the subsequent test sequence, the energy of the residual self-interference after cancellation signal $\rho[k]$, as given in (6), is measured. If $\rho[k]$ is greater than a predefined threshold, then a collision is declared.

Continue-transmission sequence. On the timeline in Fig. 5, another feature of the downlink is shown. Each time a device receives the frame header, the device decides if it wants to receive the corresponding payload. If yes, the device transmits a four-byte, random, Manchester encoded CT sequence using full-duplex transmission enabled in Sec. 3.1. The sequence is detected by the AP by measuring the energy in $\rho[k]$. When multiple devices send a random CT sequence in response to the same frame header, the random signals are added at the AP receiver, which, using energy detection, allows the AP to detect that at least one device is present. The CT sequence enables APs to abort transmitting the frame payload if no devices are present, leading to more efficient channel usage and a reduction in power consumption at APs.

Uplink. As shown in Fig. 5, the devices are only allowed to transmit uplink frames when they are receiving a downlink frame. This eliminates the possibility that the devices located at neighbouring APs can interfere. However, due to the limited transmission power of the infrared LED at each device, the NLOS link (via ceiling reflection) cannot be used for CCA or CD, like the downlink protocol, to prevent uplink intra-AP interference. Therefore, the devices are hidden from each other. An uplink contention-based protocol regulates the transmission of uplink frames during the downlink frame. In FD-CSMA/CD-HA, the hidden-device problem is solved by interleaving 'busy slots' at $80\mu s$ intervals with frame payload data during downlink transmissions, as illustrated in Fig. 6. A busy slot is either 'LOW-HIGH' to indicate that the uplink is idle or 'HIGH-LOW' to indicate that it is busy. Devices decode the latest received busy slot to determine the uplink channel status. Hence, the busy slots solve the hidden device problem in a similar way to the dedicated busy tone channel in BTMA [5]. We refer to our mechanism as *busy slot multiple access*.

The advantage of busy slot multiple access over BTMA is the reduced hardware complexity since, unlike busy tone multiple access, no secondary frequency channel is required.

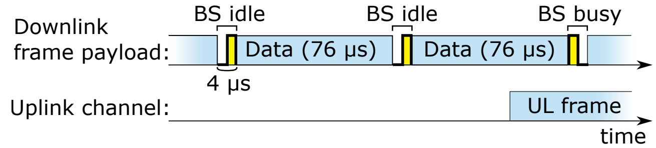

Figure 6. The proposed busy slot mechanism to solve the hidden-device problem

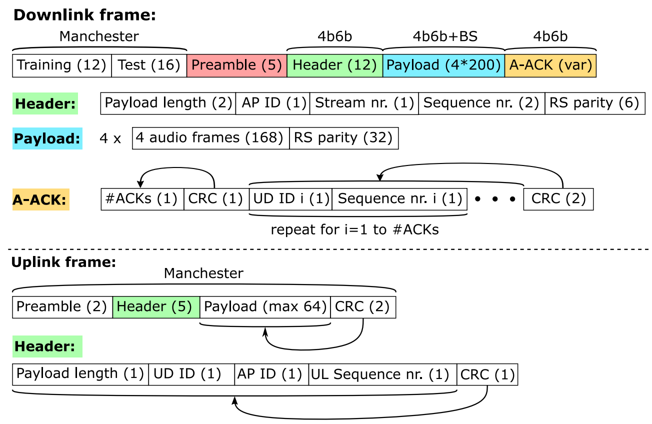

Figure 7. The downlink and uplink frame formats in LightTour. The field lengths are expressed in number of bytes (A-ACK: aggregated acknowledgement; RS: Reed-Solomon; UD: User Device.)

The disadvantages of busy slot multiple access are twofold. Firstly, there is a reduction of downlink physical layer data rate by $\frac{4\mu s}{80\mu s} = 5\%$. Secondly, due to the non-zero time interval between successive busy slots, the vulnerable period (time during which the devices can start transmitting simultaneously) of the uplink protocol is increased, resulting in an increased collision probability. However, the resulting uplink throughput reduction remains limited ($< 10\%$) since the average vulnerable time period increase (80 μs) is only 2% of the uplink frame length [24, page 510]. The busy slot interval (80 μs) can be tuned to control the trade-off between downlink and uplink throughput.

Not highlighted in Fig. 5 but will be illustrated in Fig. 7 is the short aggregated-acknowledgement message transmitted by the AP at the end of each downlink frame. The aggregated-acknowledgement acknowledges all successful uplink transmissions that occurred during this downlink frame. By doing this, devices can update their backoff exponents and retransmit any unsuccessful frames at a later time.

Frame format. The downlink and uplink frame structures are shown in Fig. 7. A downlink frame starts with a random training-test sequence. The training part is used to estimate $\hat{g}$ and $\hat{d}$, as presented in Sec. 3.1. The test part is used to detect collisions. Next, the frame contains a short preamble to facilitate frame detection and clock recovery. The header is unified, containing MAC and application layer related fields. The payload has four (200,168) Reed-Solomon blocks. Each block carries four Speex [37] audio pieces, each piece lasting for 20 ms. Hence, each downlink frame carries an audio pieces of $4 \times 4 \times 20$ ms $= 320$ ms. We encode the frame header with a short (12,6) Reed-Solomon code, allowing devices to quickly decode the received headers and immediately respond with a CT sequence. An aggregated acknowledgement is appended to the end of the downlink frame and contains acknowledgements for all successful uplink frame transmissions during the current downlink frame.

An uplink frame is much shorter compared to a downlink frame. The uplink frame contains a 2-byte preamble, a unified header with a 1-byte cyclic redundancy check (CRC) and a payload of up to 64 bytes with a 2-byte payload CRC. For the uplink, CRC is sufficient because uplink frame retransmissions are possible and thus some uplink frame error rate is acceptable. Thus, the more complex Reed-Solomon code is not necessary (CRC requires 98% less computation time compared to Reed-Solomon codes on our prototype).

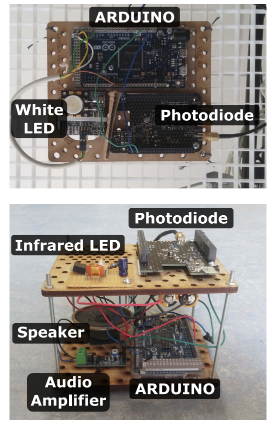

Figure 8. Top: AP; Bottom: (user) device

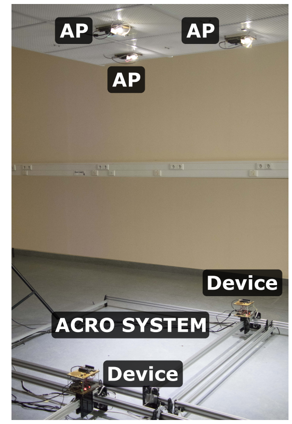

Figure 9. Experimental setup

3.3 Application Layer

The AP application layer broadcasts audio streams to devices. Each AP stores a different audio track to provide a location-based service. An audio stream is a series of chronological audio frames corresponding to this AP's audio track. A leaky bucket algorithm [29, page 407] is adopted to provide both flow control and congestion control. Using the leaky bucket, each device has a small buffer corresponding to 4*s* of audio data. The first few frames of a stream are transmitted quickly until the buffer (i.e. bucket) is full. Subsequent audio frames are transmitted at the same interval as they are played back at the device. From our experiments, the congestion control provided by the leaky bucket mechanism also significantly increases the downlink fairness under high loads.

By default, an AP periodically (period of 1*s*) broadcasts short probing frames. The payload of a probing frame corresponds to the start of that AP's audio track. Probing frame transmission is aborted when no CT sequence is detected. When a device wants to listen to the audio stream corresponding to a probing frame, it sends a CT sequence. The CT sequence is sent for each subsequent frame belonging to that stream. Transmission of further audio frames of a stream is aborted if no CT sequence is detected by the AP for two subsequent frames of a stream.

A device tracks the RSSI of all frames (including probing frames) it receives from all nearby APs. The device only listens to one stream from the AP with the highest RSSI. To avoid frequent switching when the device is at the edge between APs, a hysteresis mechanism similar to cellular systems is employed [10].

The AP broadcasts audio streams to the devices in its coverage. That is, multiple devices can listen to the same audio stream. This allows each AP to serve a large number of devices with a low PHY data rate. When sufficient channel capacity is available, an AP can transmit multiple (up to three in our prototype) audio streams in parallel, each of a different part of the audio track. If an AP is transmitting a stream and enough channel capacity is available, it additionally broadcasts probing frames to allow arriving users to open a new stream. Parallel streams hence allow a user, when arriving at an AP, to quickly receive the start of an audio stream, even when other devices are already present and listening to a stream.

On the uplink, the device sends visitors' feedback on the museum exhibition to the nearest AP. The feedback can either be short text messages or audio pieces.

4 System Implementation

We use off-the-shelf components to implement LightTour. Fig. 8 shows the *hardware*. Each AP/device has an Arduino Due microcontroller [4]. We use the LED transmitter front-end (with a 2.5 Watt, 150 lumen CREE XT-E LED) and photodiode receiver front-end (with a S5971 photodiode) that are available in the community [6, 7]. The operating modes of the LED transmitter front-end are explained in Sec. 2. The receiver front-end consists of a photodiode, amplifier, bandpass filter and 1MHz ADC. For the device, we build our own infrared LED front-end consisting of an NMOS driver circuit and infrared LED (TSHG5510 [32]) with an optical output power of 75 mW. Each device is also equipped with a small speaker and audio amplifier for demonstration and validation of the protocols.

We develop the *software* of the AP and the device in C from scratch but we do use well-known libraries such as the Speex audio codec [37]. Since most functions require real-time processing we use techniques such as direct memory access, interrupt-based scheduling and -O3 compilation to minimize the execution time. The symbol rate of our system is 500 ksymbols/s, which corresponds to a symbol time of 2 μs.

For the downlink, continuous clock recovery based on 2-fold blind oversampling is implemented in the software for downlink reception at the devices [18]. Clock recovery is necessary because of the microcontroller clock drift in combination with the long downlink frame length (27 ms). For the uplink, the frame length is much shorter (3.2 ms) and therefore, clock drift mitigation is not required.

5 Performance Evaluation

In this section we evaluate the system performance of LightTour. Our experimental setup includes a network of three APs and two devices. Fig. 9 shows a snapshot of the

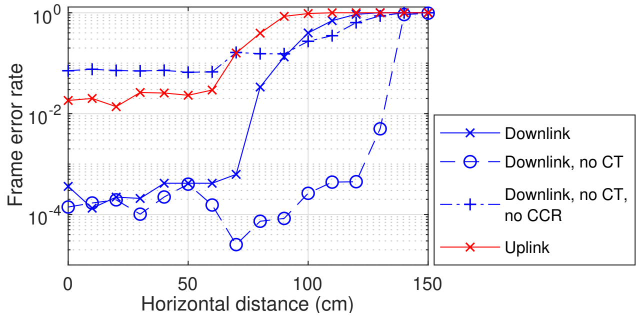

Figure 10. Frame error rate vs horizontal AP-device distance (CCR: continuous clock recovery)

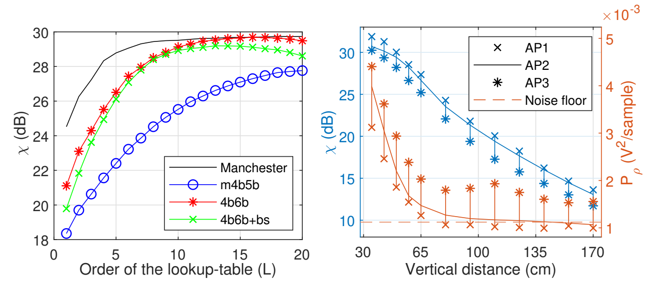

Figure 11. Performance of the proposed SIC mechanism

setup. The APs are mounted at a vertical distance of 1.7 m from the floor, forming an equilateral triangle with a side length of 1.2 m (a motivation for these distances is given in Sec. 7). The two devices are placed on the floor at a height of 0.1 m, each controlled by an OpenBuilds ACRO system [23]. The devices can be moved to any position within an area of 1.5 m x 1.5 m.

5.1 Point-to-point Link

Data rate. For the point-to-point link with a single AP communicating with a single device, a MAC layer data rate of 190 kb/s in the downlink and 71 kb/s in the uplink are achieved in our experiments. As such, LightTour improves the downlink data rate by 20 times compared to the state-of-the-art system in [26] after normalizing to the CPU clock speed, as detailed in Table 1. The achieved data rate in LightTour is also comparable to the latest OpenVLC1.3 after normalizing to the CPU speed. Furthermore, OpenVLC1.3 does not support full-duplex communication [12].

Frame error rate. We evaluate the point-to-point frame error rate versus the AP-device horizontal distance at a AP-device vertical distance of 1.7 m. The results are shown in Fig. 10. On the downlink, failed frames are not retransmitted. Hence, an increase in the frame error rate on the downlink leads to dropped audio frames and thus a decrease in audio quality. Empirically, it was found that a downlink frame error rate above 5% makes the audio stream difficult to understand. As a result, the useful horizontal transmission range equals 80 cm with the CT mechanism and 130 cm without it. Enabling continuous clock recovery (Sec. 4) on the downlink lowers the frame error rate from 7% to 0.05% at an AP-device distance of less than 1 m. Hence, the clock recovery is essential for a satisfying listening experience. The uplink has the ability to retransmit the same frame if an error occurs. Thus, an increase in the frame error rate for the uplink results in more retransmissions which lowers the effective throughput. The best uplink performance is obtained at AP-device distances less than 0.6 m. At high SNR, achieved for small AP-device horizontal distance, the uplink frame error rate equals 2%.

5.2 Self-Interference Cancellation

In this section, we evaluate the performance of the proposed SIC mechanism on our prototype. First, we study the effect of the line code and the lookup-table order L. We place the AP pointing downwards at a vertical distance of 65 cm from the floor, and perform SIC at the AP. We define the metric $\chi = P_s/P_\rho$ to quantify the SIC performance, where P_s is the received self-interference power before cancellation and P_ρ is the residual self-interference power after cancellation.

In the experiment, we measure χ for various values of L and for different line codes at the same P_s. The result is shown in Fig. 11(left). As expected, we observe that a larger filter order leads to an increase in χ. For example, when the lookup-table order is greater than 10, nearly 30 dB cancellation is achieved with Manchester coding. The SIC performance for non-DC-balanced line codes (e.g., m4b5b [17]) is significantly worse than DC-balanced line codes. This is because non-DC-balanced signals have a non-zero power spectral density at low frequencies which are effectively removed by the receiver front-end bandpass filter. Therefore, to achieve a high χ, the SIC filter $\boldsymbol{h_r}$ (approximated by the lookup-table $\mathcal{T}$ in (3)) needs to attenuate these low frequency components to the same extent as the hardware bandpass filter. This is impossible with a low $\boldsymbol{h_r}$ filter order (i.e., low lookup-table order).

Based on the above results, we use a filter order of $L = 11$, corresponding to $2^{11} = 2048$ kB of memory required to store the lookup-table, and the 4b6b line code in the rest of our experiments. This setting strikes a balance between the amount of cancellation, memory requirement and line code efficiency.

The performance of the SIC algorithm in terms of different AP-floor distances is given in Fig. 11(right). The power P_ρ and other power signals in the rest of this work are measured in terms of the voltage at the ADC output, with unit V^2/sample. Varying the AP-floor distance allows simulating scenarios with various values of P_s. At AP-floor distances larger than 70 cm, the SIC achieves P_ρ close to the noise floor indicating that the SIC is sufficiently independent on the path loss h_c. The ratio χ decreases at large distances, because P_s decreases while P_ρ is lower-bounded by the noise floor. As will be shown in Sec. 6, the SIC performance plays an important role in the whole network performance.

The execution times per received sample of the SIC algorithms based on least-mean-squares, finite-impulse-response and lookup-table for $L = 11$ are shown in Table 2. The symbol time is 2 μs (cf. Sec. 4). Additional operations at the AP such as uplink frame decoding and preamble detection require at least 1.1 μs per sample. As a result, real-time SIC is not possible with least-mean-squares nor finite-impulse-response but is possible with our lookup-table based method.

Table 1. **Comparison of LightTour's MAC layer data rate with state-of-the-art solutions**

	Platform	CPU Speed	Downlink data rate	Uplink data rate
IoTDI'17 [26]	ARM M0	48 MHz	$\approx$ 5 kb/s	not specified
OpenVLC1.3 [12]	BeagleBone Black	2 x 200 MHz	$\approx$ 450 kb/s	not supported
LightTour	Arduino Due	84 MHz	190 kb/s	71 kb/s

Table 2. Comparison of SIC algorithms ($L = 11$)

Algorithm	Execution time (μs)
Least-mean-squares based SIC	3.40
Finite-impulse-response based SIC	1.56
Our lookup-table based SIC	0.43

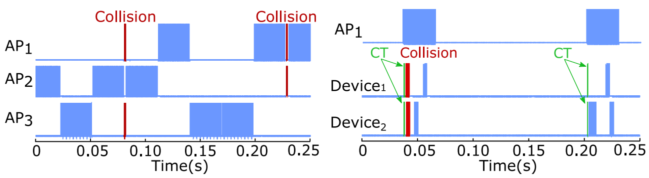

Figure 12. Oscilloscope snapshot of FD-CSMA/CD-HA

5.3 MAC protocol

We use our prototype setup in Fig. 9 for our MAC protocol evaluations.

Snapshot. An oscilloscope snapshot of our running MAC protocol is given in Fig. 12. We first consider a saturated scenario in which all three APs always try to send data to the devices when the downlink channel is free. We observe from Fig. 12(left) that, thanks to the full-duplex transmission, the APs can detect collisions perfectly. The APs stop their transmissions immediately once a collision is detected, leading to an increase in channel utilization. Fig. 12(right) shows the downlink and uplink transmissions in an unsaturated scenario. It is clear that the hidden-device problem is alleviated by our proposed busy slot multiple access. Note that uplink collisions are still possible (as highlighted in red in Fig. 12(right)) when two or more devices start transmitting during the vulnerable time period caused partly by the non-zero busy slot interval presented in Sec. 3.2.

Downlink throughput. We evaluate the downlink channel utilization versus the offered load for our FD-CSMA/CD HA protocol without considering the application layer. Using this technique, a general performance comparison between FD-CSMA/CD HA and CSMA/CA is possible. Following general testing methodology for CSMA protocols [28], the frame arrival at each AP follows a Poisson distribution with arrival rate λ. The offered load $G = \lambda M$ is the total average number of frame arrivals per frame time for a network of M APs [33]. We use $M = 3$ and $N_{cca} = 80$ samples as the default value of the window length for CCA. Fig. 13(left) shows the results. Our proposed FD-CSMA/CD-HA MAC protocol achieves a channel utilization of 0.958 under high load. Additionally, the three APs have approximately the same channel utilization, indicating that the protocol ensures fairness among the APs.

To evaluate the performance of the downlink MAC protocol of LightTour in a large-scale scenario ($M > 3$), we build

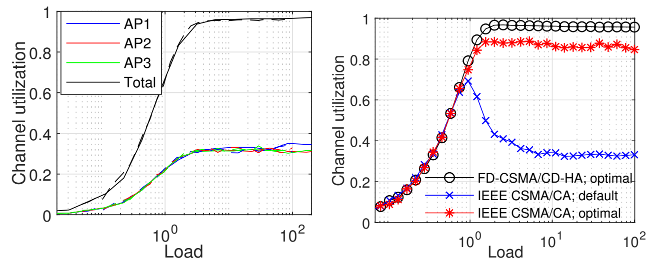

Figure 13. Left) Downlink MAC layer throughput vs offered load G (solid line: *simulation*, dashed line: experiment); Right) *simulation* results for a large-scale scenario with $M = 9$ APs

Table 3. Optimal MAC parameters from simulations

Protocol	p	macMinBE	macMaxBE	macMaxNB
IEEE CSMA/CA default	0	3	5	4
IEEE CSMA/CA optimal	0	4	9	5
FD-CSMA/CD-HA optimal	0	0	4	9

a Monte-Carlo simulator in Matlab. We assume perfect SIC, CCA and CD and use the same protocol parameters (e.g., frame length, propagation delay, MAC parameters) as in our experimental setup. As a validation, we first compare the simulation results for three APs with the experimental results. The comparison is given in Fig. 13(left), showing that the simulation matches with the experiment.

We use our simulator to study a large-scale scenario with nine interfering APs (cf. Fig. 1). We compare our solution with IEEE 802.15.7 CSMA/CA both for default values of the 802.15.7 [1] MAC parameters (*p, macMinBE, macMaxBE, macMaxCSMABackoffs*) and for optimized values of these parameters. For a description of these parameters we refer to the IEEE 802.15.7 standard document [1]. Optimal values are found by a brute force search using our simulator over the parameters, optimizing a multi-objective cost function consisting of the sum of three normalized performance metrics: total channel utilization, average frame delay and Jain's fairness index [9] over a wide range of traffic loads. The default and optimal MAC parameters are shown in Table 3. For both CSMA/CA and FD-CSMA/CD HA, non-persistent sensing (p=0) is optimal. The throughput results are shown in Fig. 13(right). We can see that under high traffic load, our method improves channel utilization by 185% compared to the default IEEE CSMA/CA protocol and 10% compared to the IEEE CSMA/CA protocol with optimal MAC parameters. The optimal FD-CSMA/CD HA parameters are used in all experiments involving the downlink MAC protocol.

Finally, we add the application layer to our simulator to evaluate the performance of our MAC protocol for the LightTour application in a large-scale setup. We simulate a room

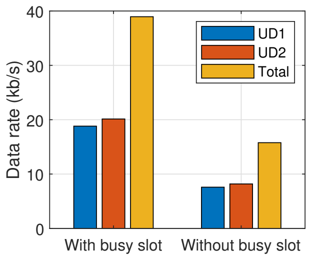

Figure 14. Uplink data rate with and without busy slot (UD: user device)

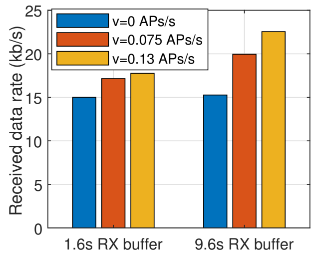

Figure 15. Received downlink data rate (v: speed of the user device)

consisting of 6×6 APs (similar to Fig. 1) with 1.2 m horizontal distance between APs, a horizontal AP cell size of 2 m and up to 100 mobile devices. Each AP can sense only its immediate neighbours (incl. diagonal neighbours). Each device follows a random walk. Our simulator throughput results indicate that congestion control at the application layer with the leaky-bucket algorithm is necessary to achieve fairness between APs for this grid setup. Without the leaky-bucket algorithm, APs at the center of the grid (competing with 8 other APs) have a significantly lower throughput than APs at the edges (competing with only 3 to 5 other APs). With the leaky-bucket algorithm, we achieve a fair channel utilization of 0.2 per AP (each AP uses on average 20% of the available channel bandwidth). For comparison, TDMA in a square grid of APs can achieve a theoretical channel utilization of 0.25 per AP, since a spatial reuse factor of 1/4 ensures that no two neighbouring APs can transmit at the same time. However, TDMA require AP synchronization and coordination, while our system works in a fully distributed manner.

Uplink throughput. We measure the throughput for two saturated devices transmitting to an AP. Fig. 14 shows the results when the busy slot mechanism is enabled and disabled. The busy slot mechanism improves the total uplink throughput from 15.8 kb/s to 39 kb/s and achieves a throughput efficiency of $\frac{39\text{kb/s}}{71\text{kb/s}} = 55\%$ in which 71 kb/s is the uplink throughput with only a single device (cf. Sec. 5.1).

6 Robustness evaluation

Lastly, we study the factors that might reduce the system performance, including device mobility, ambient light interference, blockage, and imperfect CCA and CD.

6.1 Mobility

To analyze the effect of the mobility of devices on the application layer, we use the ACRO system to continuously move one of the two devices between the three APs while the other device is stationary. The received data rate at each device is measured for several values of the device receiver buffer size and device movement speed. We define the device movement speed, denoted by v, as the average number of APs visited per second by that device.

The measurement results are shown in Fig. 15. We observe that an increase in device mobility leads to an increase in the average received data rate. This is because the leaky bucket algorithm at each AP transmits the first few frames of each audio stream at a higher rate compared to later frames.

Table 4. Various metrics versus ambient light level (DL FER: downlink frame error rate)

Ambient light (lux)	RX noise (V^2/sample)	DL FER	SIC χ (dB)
15	1.378	10^{-4}	29.91
500	2.8	10^{-4}	26.92
1000	5.37	6×10^{-4}	24.27
1800	11.36	5×10^{-4}	23.65

Every time a device moves to a new AP, the device receives the first few frames of the new AP's audio streams that are transmitted at this higher rate. Similarly, a larger device receiving buffer size leads to an increase in receiving data rate since the fraction of frames that are transmitted at a higher data rate is larger. In conclusion, our experiment shows that device mobility has no negative impact on the downlink throughput.

6.2 Ambient light

We investigate the effect of ambient light on the system performance. First, we analyze the effect on the following metrics: receiver noise level, downlink frame error rate and SIC performance. To create different ambient light levels, a transmitter front-end operating in illumination mode is placed at different distances from the receiver. The measurement results are shown in Table 4. As a reference, to preserve artworks the recommended art museum light levels should not exceed 300 lux [11]. Normal office light levels range from 500 lux to 1000 lux [21].

We observe that more ambient light leads to more receiver noise, which is measured at the ADC output. This is expected, as ambient light causes shot noise on the photodiode [14]. Ambient light impacts the SIC performance (χ), which is not only reduced due to extra shot noise but also due to the receiver front-end non-linearity [38].

Both receiver noise and SIC performance affect the accuracy of the CCA and CD mechanisms at the AP. This is important, as imperfect CCA and CD impacts the downlink MAC protocol performance (cf. Sec. 6.4). We use the experimental setup in Fig. 9 to measure the CCA and CD receiver operating characteristic curves for light intensities of 15 and 500 lux at the AP receiver on the ceiling (pointing downwards). The choice of these values is supported by following empirical observation: in an artificially lit room with realistic height of 2.7 m in which 500 lux is measured at a height of 1 m from the floor, 40 lux is measured at the AP receiver on the ceiling. We analyze the impact of the CCA window length N_{cca} (in number of samples) and CD training + test sequence length L_{cd} (in number of bytes) on the CCA and CD accuracy.

The measured receiver operating characteristic curves are shown in Fig. 16. Some curves, e.g., the curve for 15 lux with $N_{cca} = 40$ samples, are not visible because they result in false positives and false negatives rates below 10^{-6}. We make three important observations: (1) ambient light raises the false positive rate (FPR) and the false negative rate (FNR) of CCA and CD; (2) the negative impact of ambient light can be largely mitigated by an increase in N_{cca} and L_{cd}; (3) CD requires more samples to achieve the same FNR for a given FPR compared to CCA.

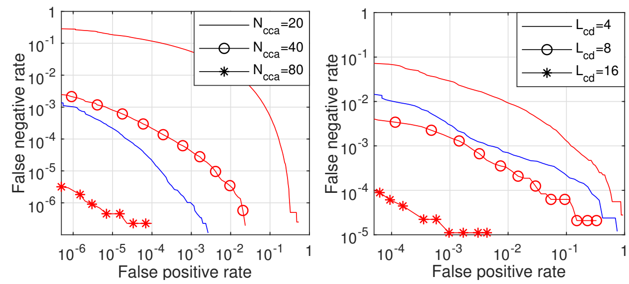

Figure 16. Receiver operating characteristic curves for CCA (left) and CD (right) at ambient light level of 15 lux (blue) and 500 lux (red)

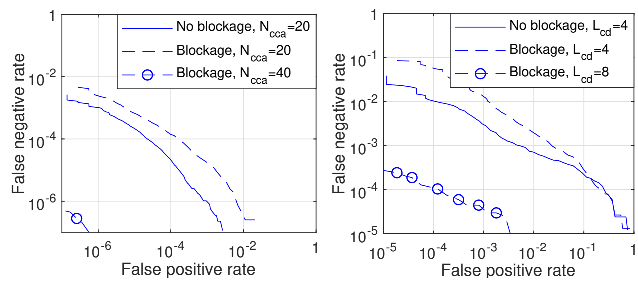

Figure 17. Receiver operating characteristic curves for CCA (left) and CD (right) with and without blockage

6.3 NLOS blockage

Blockage, such as caused by people and objects, can affect the received signal strength of the inter-AP NLOS signal. In this subsection, we study the effect of blockage on the CCA and CD receiver operating characteristic curves. Similarly to the work in [8], we model a person by a black cylinder with 40 cm diameter and a brown cardboard cover at the top. As shown in Fig. 18, we place the cylinder between two APs and set the vertical distance between the APs and the top of the cylinder to 1 m. This is equivalent to a real-world scenario with a ceiling height of 2.7 m and visitor height of 1.7 m. The blockage measurements are done at an ambient light level of 40 lux measured at the AP receiver (cf. Sec. 6.2). In this setup, the received signal strength of the inter-AP NLOS signal is reduced from $8.0 \times 10^{-3}\,V^2/\text{sample}$ (without blockage) to $5.2 \times 10^{-3}\,V^2/\text{sample}$ (with blockage). Fig. 17 shows the impact of the blockage on the CCA and CD ROC curves. We conclude that, for both CCA and CD and for a given FPR, blockage increases FNR by up to an order of magnitude. However, similarly to ambient light, the effect of NLOS blockage can be largely mitigated by increasing N_{cca} and L_{cd}.

6.4 Imperfect CCA and CD

The impact of imperfect CCA and CD on the CSMA/CD protocol has been studied to some extent in literature [27], [16]. To study their impact on FD-CSMA/CD-HA in large scenarios, we use our Matlab simulator with nine APs presented in Sec. 5.3 and simulate the FPR/FNR by including a non-zero probability of false positive/negative each time CCA and CD is performed.

False positive rate (FPR). Non-zero CCA and CD FPR results in an increase in the delay between frame arrival in the transmission buffer and actual start of frame transmission. Therefore, this impacts the throughput. Fig. 19(left) shows the throughput curves under high load with CCA and CD FPRs of 0 and 0.2. We notice that values up to 0.2 of the CCA and CD FPRs have an impact of maximum 3% on the throughput. A surprising result is that small non-zero values of CCA FPR improve the throughput at high loads. This is because the optimal MAC parameters are calculated using an objective function that incorporates throughput, delay and fairness over a wide range of traffic loads (cf. Sec. 5.3). Higher throughput at high loads can be achieved at the cost of other performance metrics such as the frame delay. A non-zero CCA FPR decreases the mean transmission probability which leads to a slightly larger frame delay at low loads, as shown in Fig. 19(right). Overall, we conclude that FPR has little impact on the MAC performance.

False negative rate (FNR). A non-zero CCA and CD FNR results in frame collisions and thus an increase of the FER. We study the worst-case scenario in which frames are completely lost at all device locations when at least two of the nine APs transmit part of their frames simultaneously. In reality, the signal-to-interference-plus-noise ratio (SINR) (interference due to nearby AP transmission) at most device locations is sufficiently high such that simultaneous transmission of two APs does not always result in a lost frame at the device. Fig. 20 shows the simulation results at a load $G = 20$ for CCA (with CD FNR=0) and CD (with CCA FNR=0). The vertical axis denotes the probability that a frame collides with another transmission because of a false negative CCA or CD. We notice that a CCA FNR of 3×10^{-6} leads to a collision probability of 10^{-3}. For the CD, the same collision probability is achieved for a FNR of 3×10^{-3}.

We conclude that blockage and ambient light can negatively impact the performance of our protocol. However, by properly selecting the parameters such as the CCA and CD energy thresholds, N_{cca}, and L_{cd}, these negative impacts can be largely reduced and even eliminated.

7 Discussions

In this section, we present the current limitations of LightTour. We also discuss some potential applications.

Prototype dimensions. In our prototype, APs are placed at a height of 1.7 m with 1.2 m inter-AP horizontal spacing. At these distances the SNR of the inter-AP NLOS signal is measured to be 8 dB as shown in Fig. 2. Sufficient inter-AP SNR is necessary to achieve accurate CCA and CD. We now argue that in a real-world museum setup, with increased ceiling height and greater inter-AP spacing, inter-AP SNR similar to our prototype can be achieved by tuning certain system parameters. Firstly, the luminous flux of our prototype's white LEDs are 150 lumen. The DC gain of the first reflection in a visible light channel is proportional to $\frac{1}{D_1^2 D_2^2}$ [13], with D_1 the distance between an LED and a reflective point on the floor and D_2 the distance between the reflective point and receiver. Assuming a real-world, angle-preserving setup with all distances from our experimental setup scaled up by 60% (ceiling height of 2.7 m and inter-AP distance of

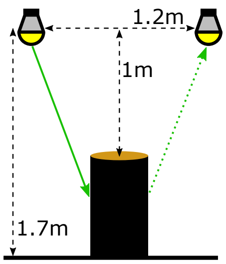

Figure 18. Experimental setup for studying the effect of blockage

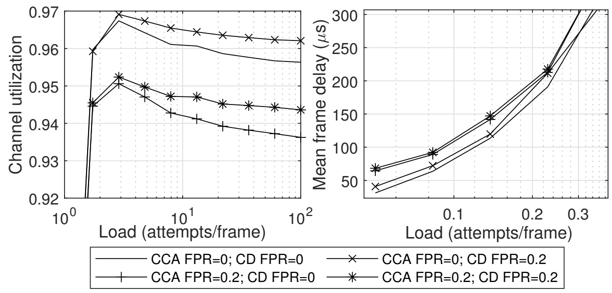

Figure 19. Result of the FPR of CCA and CD: left) downlink MAC throughput vs offered load; Right) downlink MAC mean frame delay vs offered load

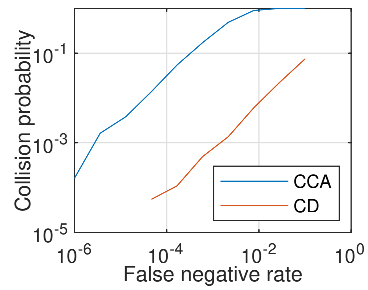

Figure 20. Worst-case collision probability versus CCA and CD FNR for load $G = 20$

1.9 m), the required LED luminous flux to achieve the same inter-AP SNR as our prototype is:

$$\frac{(2.7m)^2(2.7m)^2}{(1.7m)^2(1.7m)^2} \times 150\,\text{lumen} = 954\,\text{lumen},$$

which is in the range of regular LEDs for indoor illumination. A luminance of 954 lumen results in an average room illumination of:

$$\frac{954\,\text{lumen}}{(1.9m)^2} = 272\,\text{lux},$$

which is less than the recommended maximum illumination for artwork preservation of 300 lux [11]. Other system parameters that can be tuned to improve inter-AP SNR are floor reflectivity, AP horizontal spacing and LED half power semi-angle. Finally, if the inter-AP SNR is still insufficient, increasing N_{cca} and L_{cd} significantly decreases FPR and FNR for CCA and CCD at the cost of limited ($< 5\%$) throughput reduction due to the extra overhead. The experiment in Fig. 16 demonstrates the achieved decrease in FPR and FNR for CCA and CD for multiple values of N_{cca} and L_{cd}. Additional care must be taken when selecting abovementioned system parameters since optimal MAC performance is obtained when each AP only interferes with its immediate neighbors.

Look-up table training. In the current implementation, the look-up table $\mathcal{T}$ of each AP is learned offline and hence immutable. SIC close to the noise floor is achieved under various self-interference signal strengths (Fig. 11 and ambient light levels (Table 4). To increase practicality and SIC robustness, in future work, $\mathcal{T}$ can be trained online on each microcontroller using periodically transmitted training frames with least-square channel estimation. Factors such as temperature, circuit aging and significant shifts in ambient light can affect the RX circuit response and hence $\mathcal{T}$. Since these factors vary slowly, the interval between training frames can be large (e.g., 2 minutes), such that the long computation time for least-squares channel estimation (3.4 μs/symbol) is irrelevant.

Data rate and audio quality. Although the audio quality of 16.8 kb/s wideband speech coded using the Speex codec is acceptable, a museum might also be interested to incorporate high quality music into the audio streams. For example, the recommended bitrate for streaming stereo music with the Opus codec is 64 kb/s [15]. With three parallel 64 kb/s streams and a channel utilization of 0.2 per AP, a data rate of 960 kb/s per AP is required which is five times higher than the achieved 190 kb/s in this work. In our future work, PHY data rate can be increased by decreasing the symbol length (currently 2 μs) and/or using one of the higher-throughput modulation schemes described in the IEEE 802.15.7 standard [1] (e.g., pulse amplitude modulation). Handling faster data rates will require an adaptation of the SIC mechanism to limit the excessive table size, a redesign of the TX/RX front-ends and faster microcontrollers.

Protocol applicability. The PHY and MAC layer mechanisms proposed in this paper solve certain universal problems that harm VLC systems. Firstly, the SIC mechanism enables full-duplex transmission with only a single optical antenna and enables the CD mechanism needed for CSMA/CD. Secondly, the proposed CSMA/CD protocol is novel for VLC and allows reliable broadcasting with many uncoordinated APs, a requirement that cannot be met by IEEE CSMA/CA. Additionally, our experiments show that CSMA/CD improves (both unicast and broadcast) MAC-layer throughput under high load by 10% over IEEE CSMA/CA. Thirdly, the busy slot mechanism provides a simple solution to the hidden device problem for uplink transmissions, a common problem in VLC. Thus, we argue that our proposed solutions can be adopted individually, or collectively in other applications as discussed below.

Examples of similar applications. The LightTour protocol stack is specifically designed for a dense layout of APs that broadcast location-based information to a large number of devices and can be adopted to similar applications. For example, in airports or train stations departure times and announcements can be broadcast. Similarly, grocery shoppers can receive location-based information about products. An indoor positioning system is feasible where APs broadcast their locations together with the building's map. A potential industrial application is inventory management where warehouse shelves send and receive status updates using VLC.

8 Related work

We summarize the most relevant works in this section.

Self-interference cancellation (SIC) has been well-studied in RF communications. For example, [2] introduces an all-digital SIC mechanism for RF networks, where the channel is estimated using a least square estimator with time-orthogonal (uncorrelated) training sequences at the start of each frame. Our SIC mechanism is partly inspired by this work. Leveraging SIC for CD has also been extensively studied [31, 30]. *Compared to these state-of-the-art research on RF, our SIC mechanism for optical wireless networks has a low complexity and can be easily implemented with low-end platforms such as the Arduino Due used in this paper.*

The performance of full duplex VLC is studied in [39]. It introduces two full-duplex contention protocols: U-ALOHA and full-duplex CSMA. They state that in visible light communication environments self-interference can be ignored because of the lower received power from NLOS links compared to LOS links. *However, in our LightTour system we have showed that the self-interference from the floor reflection cannot be simply ignored.* The authors in [35, 34] develop a CSMA/CD protocol with hidden avoidance for LED-to-LED communication. Bidirectional communication is achieved using an OOK modulation scheme where the uplink channel is sensed during the transmission of LOW symbols. Furthermore, the bidirectional link enables CD. *Compared to this work, our system supports full-duplex (not only bidirectional) transmissions. Moreover, we study the CD accuracy in a much larger range of SNR conditions.*

9 Conclusion

We have designed, implemented, and evaluated a networked system with light streaming location-dependent audio to users. The proposed FD-CSMA/CD-HA protocol is enabled by a novel low-complexity SIC mechanism which provides real-time SIC on a low-cost microcontroller. Our results validate the feasibility and show the advantages of our system. As we advance, we envision that our system can inspire more work on practical full-duplex communication with light and more application scenarios.

Acknowledgment

This work is supported in part by Research Foundation Flanders SB PhD fellowship under grant number 1S59918N.

10 References

[1] IEEE Standard for Local and metropolitan area networks-Part 15.7: Short-Range Optical Wireless Communications. *IEEE Std 802.15.7-2018 (Revision of IEEE Std 802.15.7-2011)*, 2019.

[2] E. Ahmed and A. M. Eltawil. All-digital self-interference cancellation technique for full-duplex systems. *IEEE Transactions on Wireless Communications*, 2015.

[3] O. Alsulami, M. Alresheedi, and J. Elmirghani. Infrared uplink design for visible light communication systems with beam steering, 2019.

[4] Arduino. Arduino Due. https://store.arduino.cc/due.

[5] D. Barrett and T. King. *Computer Networking Illuminated*. 2005.

[6] J. Beysens, A. Galisteo, Q. Wang, D. Juara, D. Giustiniano, and S. Pollin. DenseVLC: A Cell-free Massive MIMO System with Distributed LEDs. In *ACM CoNEXT*, 2018.

[7] J. Beysens, Q. Wang, A. Galisteo, D. Giustiniano, and S. Pollin. A Cell-Free Networking System With Visible Light. *IEEE/ACM Transactions on Networking*, 2020.

[8] J. Beysens, Q. Wang, and S. Pollin. Improving Blockage Robustness in VLC Networks. In *COMSNETS*, 2019.

[9] D.-M. Chiu and R. Jain. Analysis of the increase and decrease algorithms for congestion avoidance in computer networks. *Comput. Netw. ISDN Syst.*, 1989.

[10] N. Ekiz, T. Salih, and et al. An overview of handoff techniques in cellular networks. *Information Technology*, 2005.

[11] Feilo Sylvania. *Lighting for Museums and Galleries*, 2015.

[12] A. Galisteo, D. Juara, and D. Giustiniano. Research in visible light communication systems with openvlc1.3. In *IEEE WF-IoT*, 2019.

[13] F. Gfeller and U. Bapst. Wireless in-house data communication via diffuse infrared radiation. 1979.

[14] G. Ghassemlooy, W. Popoola, and S. Rajbhandari. *Optical Wireless Communications System and Channel Modelling in MATLAB*. Taylor & Francis Group, 2013.

[15] Hydrogenaudio. Opus. `https://wiki.hydrogenaud.io/index.php?title=Opus#Bitrate_performance`.

[16] M. Kaneko. Throughput Analysis of CSMA With Imperfect Collision Detection in Full Duplex-Enabled WLAN. *IEEE Wireless Communications Letters*, 2017.

[17] T. G. Kang, S.-K. Lim, D. H. Kim, I. S. Jang, and D. W. Han. Ieee 802.15.7 vlc phy/mac specification proposal, Oct 2009.

[18] J. Kim and D.-K. Jeong. Multi-gigabit-rate clock and data recovery based on blind oversampling. *IEEE Communications Magazine*, 2003.

[19] G. Li, E. Geng, Z. Ye, Y. Xu, J. Lin, and Y. Pang. Indoor positioning algorithm based on the improved rssi distance model. *Sensors*, 2018.

[20] C. Medina, M. Nuñez, and et al. Led based visible light communication: Technology, applications and challenges – a survey. *International Journal of Advances in Engineering and Technology*, 2015.

[21] National Optical Astronomy Observatory. *Recommended Light Levels (Illuminance) for Outdoor and Indoor Venues*.

[22] A. F. Oncel. *On self-interference cancellation in wireless full duplex communications*. PhD thesis, Middle East Technical University, 2016.

[23] OpenBuilds. ACRO. https://openbuilds.com/builds/openbuilds-acro-system.5416/.

[24] K. Pahlavan and A. Levesque. *Wireless Information Networks*. Wiley Series in Telecommunications and Signal Processing. Wiley, 2005.

[25] Renesas Electronics Corporation. *Eye Safety for Proximity Sensing Using Infrared Light-emitting Diodes*, 2016.

[26] S. Schmid, B. von Deschwanden, S. Mangold, and T. R. Gross. Adaptive software-defined visible light communication networks. In *Proceedings of the IEEE/ACM IoTDI*, 2017.

[27] L. Song, Y. Liao, K. Bian, L. Song, and Z. Han. Cross-layer protocol design for csma/cd in full-duplex wifi networks. *IEEE Communications Letters*, 2016.

[28] H. Takagi and L. Kleinrock. Throughput analysis for persistent csma systems. *IEEE Transactions on Communications*, 1985.

[29] A. S. Tanenbaum. *Computer Networks Fifth Edition*. Pearson Education, 2011.

[30] T. Vermeulen. *Reliable and Low Delay Communication for Wireless Sensor Networks using In-Band Full Duplex*. PhD thesis, 2017.

[31] T. Vermeulen, M. Laghate, G. Hattab, D. Cabric, and S. Pollin. Towards instantaneous collision and interference detection using in-band full duplex. In *IEEE INFOCOM*, 2017.

[32] Vishay Semicondutors. *High Speed Infrared Emitting Diode, 830 nm, GaAlAs Double Hetero*, Aug 2011.

[33] K. Voulgaris, A. Gkelias, I. Ashraf, M. Dohler, and A. H. Aghvami. Throughput analysis of wireless csma/cd for a finite user population. In *IEEE Vehicular Technology Conference*, Sep 2006.

[34] Q. Wang and D. Giustiniano. Communication networks of visible light emitting diodes with intra-frame bidirectional transmission. In *ACM CoNEXT*, 2014.

[35] Q. Wang and D. Giustiniano. Intra-frame bidirectional transmission in networks of visible leds. *IEEE/ACM Trans. Networking*, 2016.

[36] Wei Zhao and K. Ramamritham. Virtual time csma protocols for hard real-time communication. *IEEE Trans. Software Engineering*, 1987.

[37] Xiph.Org. Speex: A free codec for free speech. `www.speex.org/`.

[38] K. Ying, Z. Yu, R. J. Baxley, H. Qian, G. Chang, and G. T. Zhou. Nonlinear distortion mitigation in visible light communications. *IEEE Wireless Communications*, 2015.

[39] Z. Zhang, X. Yu, L. Wu, J. Dang, and V. Li. Performance analysis of full-duplex visible light communication networks. In *ICC*, 2015.