

DenseVLC: A Cell-Free Massive MIMO System with Distributed LEDs

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

Diego Juara
IMDEA Networks Institute
Madrid, Spain
diego.juara@imdea.org

Ander Galisteo
IMDEA Networks Institute
& UC3M, Madrid, Spain
ander.galisteo@imdea.org

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

Qing Wang[†]
KU Leuven
Leuven, Belgium
qing.wang@kuleuven.be

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

ABSTRACT

LED luminaries are now deployed densely in indoor areas to provide uniform illumination. Visible Light Communication (VLC) can also benefit from this dense LED infrastructure. In this paper, we propose DenseVLC, a cell-free massive MIMO system enabled by densely distributed LEDs, that forms different beamspots to serve multiple receivers simultaneously. Given a power budget for communication, DenseVLC can optimize the system throughput by properly assigning the power budget among the distributed LEDs. We formulate an optimization problem to derive the optimal policy for the power allocation. Our insights from the optimal policies allow us to simplify DenseVLC's system design and propose a heuristic algorithm that can reduce the complexity by 99.96%. Besides, we propose a novel synchronization method using non-line-of-sight VLC to synchronize all the transmitters that will form a beamspot to serve the same receiver. We implement DenseVLC with off-the-shelf devices, solve practical challenges in the system design, and evaluate it with extensive and realistic experiments in a system of 36 transmitters and 4 receivers in an area of $3\text{ m} \times 3\text{ m}$. Our results show that DenseVLC can improve the average system throughput by 45%, or improve the average power efficiency by 2.3 times, while maintaining the requirement for uniform illumination.

CCS CONCEPTS

• **Networks** → **Wireless local area networks**; • **Computer systems organization** → **Embedded systems**; • **Software and its engineering** → **Embedded software**;

*Part of the work was done when J. Beysens was a research intern in IMDEA Networks Institute.

[†]Corresponding author.

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 '18, December 4–7, 2018, Heraklion, Greece

© 2018 Association for Computing Machinery.

ACM ISBN 978-1-4503-6080-7/18/12...\$15.00

<https://doi.org/10.1145/3281411.3281423>

KEYWORDS

Visible light communication, cell-free, massive MIMO, over-the-air synchronization, system design, implementation, evaluation

ACM Reference Format:

Jona Beysens, Ander Galisteo, Qing Wang, Diego Juara, Domenico Giustiniano, and Sofie Pollin. 2018. DenseVLC: A Cell-Free Massive MIMO System with Distributed LEDs. In *The 14th International Conference on emerging Networking EXperiments and Technologies (CoNEXT '18), December 4–7, 2018, Heraklion, Greece*. ACM, New York, NY, USA, 13 pages. <https://doi.org/10.1145/3281411.3281423>

1 INTRODUCTION

Artificial illumination consumes about 20% of the world's electricity and produces carbon emissions that are comparable to the global automobile fleet [1]. To reduce this high energy consumption, traditional fluorescent and incandescent bulbs are being replaced by energy-efficient LEDs [35]. Deploying LEDs is not only beneficial for illumination, but also enables Visible Light Communication (VLC) to piggyback on LED's illumination [8, 14]. Yet, the primary function of an LED remains illumination. To increase user comfort, uniform illumination has become an essential requirement in advanced lighting systems. However, due to the well-known Lambertian propagation property of LEDs, it is challenging to illuminate a large area uniformly and efficiently with a single LED [28]. To tackle this problem, deploying LEDs in arrays was proposed in [16] and has been investigated extensively to obtain uniform illumination [24, 31, 36].

With this dense-luminaries infrastructure, new techniques for communication have been explored to exploit the LEDs' spatial diversity by enabling the Distributed Multiple-Input-Single-Output (D-MISO) in VLC networks [7, 27, 33]. In these previous works, all the LEDs are synchronized to increase the received signal strength at the receiver. However, these works are highly energy inefficient. As saving energy is the key reason for deploying LEDs for illumination, it is essential that (1) VLC incurs limited extra power, and (2) no power is wasted. The above works do not satisfy these requirements, as all LEDs send the same data to a receiver and consume the same amount of power, even if they do not contribute equally to the performance improvement at the receiver. To improve the power efficiency in D-MISO VLC, recent works have proposed new precoding schemes to serve multiple users [29, 39].

However, these past works are purely based on simulations, only consider limited number of transmitters, static receivers and make simplistic assumptions for the synchronization among the LEDs.

Recently, researchers in Bell Labs proposed the concept of Cell-Free Massive MIMO (CFM-MIMO) in [18]. In CFM-MIMO, a large number of distributed, low-cost, and low power access point antennas are connected to a controller to serve a much smaller number of users. The proposed system is not partitioned into cells anymore and all the users are simultaneously served by all LEDs within the receiver's field of view. This facilitates mobility and improves the dynamic performance, compared to the conventional small cell-based design [13]. CFM-MIMO has now attracted increasing attentions from the massive MIMO society, and has been proved to improve greatly the system performance in terms of outage rate, system throughput, and so on [19–22].

Inspired by this new concept, in this paper we propose, design and implement *DenseVLC*, a practical CFM-MIMO system enabled by densely distributed LED luminaries. With the intrinsic characteristics of being low-cost, low power consumption, and densely distributed, DenseVLC can enable adaptive and power-efficient CFM-MIMO beamspots to serve multiple receivers simultaneously. Given the measured link qualities between the LED transmitters (TXs) and receivers (RXs), DenseVLC allocates dynamically the power budget for communication among the LEDs such that the total system throughput is maximized. DenseVLC also maintains a constant and uniform illuminance level. Our contributions are summarized as follows:

- We propose DenseVLC and design its system architecture. It enables power-efficient and adaptive CFM-MIMO with VLC to serve multiple RXs. We formulate the policy on adapting the CFM-MIMO beamspots to achieve the best system throughput as an optimization problem and derive the solution (Sec. 3).
- We show that given a constraint on power budget, using all the power budget does not necessarily mean the system will operate in the most power-efficient state. We also show that only two modes of operations for the LEDs are enough to efficiently allocate the power budget, either only illumination at the bias current or full swing in the LED communication region. The latter allows us to simplify the hardware and system design of DenseVLC (Sec. 4).
- Based on the insights obtained in Sec. 4, we further design a heuristic algorithm that can reduce the complexity of finding the optimal adaptation policy by 99.96%, at the cost of sacrificing the system throughput only by 1.8%. This allows DenseVLC to quickly react to users' dynamics (movement, etc.), which is important to maintain the optimum of the system performance in mobile scenarios (Sec. 5).
- We propose a novel and practical synchronization method by exploiting non-line-of-sight VLC to synchronize all the TXs that serve the same RX. This synchronization method is more fine-grained, scalable and flexible compared to traditional methods (Sec. 6).
- We implement DenseVLC with off-the-shelf devices. We design and develop 36 TXs and 4 RXs from scratch to build a

real CFM-MIMO system with VLC. We solve practical challenges in the system implementation, such as those related to the proposed synchronization (Sec. 7).

- We evaluate DenseVLC in different scenarios, considering limitations caused by real systems, and compare it with existing solutions. Our results show that DenseVLC improves the average system throughput by 45%, or the average power efficiency by 2.3 times, compared to existing solutions. These improvements are obtained without affecting the uniformity of illumination (Sec. 8).

2 DESIGN SPACE AND BACKGROUND

2.1 Design space

Our goal in DenseVLC is to design and develop a practical CFM-MIMO system that addresses the following challenges:

Interference. In DenseVLC, multiple RXs can be served simultaneously. Due to the existence of densely distributed LEDs in the system and mobility of RXs, severe interference can arise and degrade the performance of the system greatly. Thus, strong interference among different beamspots must be avoided.

Power efficient. DenseVLC decides which LEDs of the array form beamspots. This adaptive decision should be aware of the power consumption for communication and the SINR at RXs, and should always lead to the best system throughput within the allowed power budget for communication, without affecting the illumination uniformity.

Fast adaptation. VLC links exhibit high dynamics when the TXs and RXs are not static [38]. We are interested in mobile RXs which are more practical in reality. Therefore, algorithms are required that can quickly determine how to form the beamspots for the RXs.

Scalable and flexible synchronization. Synchronization plays an important role in practical CFM-MIMO systems because TXs need to send the same signal simultaneously. However, none of the existing related VLC systems have proposed a practical solution. As we will show in Sec. 6.1, synchronization based on traditional methods such as Network Time Protocol are neither flexible nor scalable for dense deployments. Other methods such as GPS disciplined oscillator (GPSDO) [15] cannot be applied, as denseVLC will be typically deployed indoors.

Minimal circuitry to drive LEDs. The LEDs that form CFM-MIMO beamspots are spatially distributed within the area-of-interest and should have simple electronics and logics to transmit data. This allows their integration into today's infrastructure, retrofitting off-the-shelf LEDs, and further reduce their power consumption for communication.

2.2 Background

Channel propagation. In VLC, the intensity of light beam from an LED is captured by the Lambertian propagation model. Denoting P_t as the optical power of the transmitter, the received optical power P_r at the receiver is given as:

$$P_r = H \cdot P_t + n, \quad (1)$$

where n is the noise, which is modeled as additive white Gaussian noise (AWGN), H is the line-of-sight path loss, and

$$H = \begin{cases} \frac{(m+1)A_{pd}}{2\pi d^2} \cos^m(\phi)g(\psi) \cos(\psi), & 0 \leq \psi \leq \Psi_c \\ 0, & \text{otherwise} \end{cases} \quad (2)$$

where $m = \frac{-\ln(2)}{\ln(\cos \phi_{1/2})}$ is the LED's Lambertian order in which $\phi_{1/2}$ denotes the half power semi-angle, A_{pd} the collection area of the photodiode, d the distance between the transmitter and the receiver, ϕ and ψ the irradiation angle and incidence angle, respectively, $g(\psi)$ the concentrator and filter gain of the photodiode, and Ψ_c the field-of-view of the receiver.

Operation mode of LEDs in VLC. In VLC systems, LEDs can have two operating modes: illumination mode and illumination+communication mode:

- *Illumination mode:* The LED is solely used for illumination in this mode, which can be achieved by applying a constant current to the LED.
- *Illumination+communication mode:* In this mode, the light intensity emitted from LEDs is modulated to transmit data. To avoid flickering, the average brightness of LEDs must remain the same as in *illumination mode*.

3 DENSEVLC DESIGN

We consider an indoor area with a grid of LEDs distributed on the ceiling. The LEDs illuminate the area and act as TXs to communicate with multiple RXs. The TXs are synchronized dynamically to enable CFM-MIMO beamspots.

3.1 System architecture

The system architecture we propose for DenseVLC consists of a controller, a grid of densely distributed TXs, and multiple RXs, as shown in Fig. 1.

Controller. It orchestrates the CFM-MIMO beamspots, hosts the decision logic and the MAC protocol, and guides the wireless synchronization among the TXs.

TXs. Densely distributed TXs integrate simple electronics to seamlessly switch between illumination mode and illumination+communication mode. In the latter mode, relevant TXs send synchronized signals to the RXs. This simple design of a TX facilitates its implementation in today's indoor environments. *Synchronization among the TXs is achieved through VLC*, where synchronization signal piggybacking on visible light is sent by a leading TX, reflected by the ground, then received and decoded by other TXs. This design for the synchronization makes DenseVLC scalable and flexible.

Mobile RXs. They can move within the area-of-the-interest and receive data from the TXs. A RX is able to measure the channel qualities between the TXs and itself, and send this information back to the controller for the orchestration of adaptive CFM-MIMO beamspots. Note that DenseVLC only requires channel measurements, and does not need to know the location of the RXs.

3.2 MAC protocol

The MAC protocol that can exploit the dynamic link qualities between the TXs and the RXs, due to the RXs' mobility or environment changes, works as follows:




Figure 1: DenseVLC's system architecture

- The controller sends pilot signals in a time-division scheme to each LED of the array. Each RX measures the downlink channel qualities between each TX and itself, and reports these measurements to the controller.
- Based on the measured downlink channel qualities and the allowed power budget for communication, the controller implements a decision logic to allocate the communication power among the TXs, with the goal to optimize the system performance. Based on this allocation, TXs form CFM-MIMO beamspots to serve multiple RXs.
- For each CFM-MIMO beamspot, a TX is appointed by the controller as the leading TX that sends synchronization signals through VLC to synchronize all TXs that form the same CFM-MIMO beamspot. The visible light carrying the synchronization signal is reflected by the ground and then detected by the TXs to perform synchronization.
- Synchronized TXs switch their LEDs to operate in illumination+communication mode (cf. Sec. 2.2), send synchronized signals to the desired RX. If a TX is not assigned any communication power, its LED operates in asynchronous illumination mode to only provide illumination.

Next, we introduce our system model that will be used to derive the controller's policy on allocating the communication power among the TXs that form CFM-MIMO beamspots.

3.3 System model

We consider an area with N LED transmitters deployed in a grid on the ceiling. For simplicity, we assume that each TX uses a single LED¹. Each LED has two operating modes, as introduced in Sec. 2.2, and can switch seamlessly among them. When an LED is in the illumination mode, a constant current I_b flows through the LED to achieve the desired illumination. When it operates in the illumination+communication mode, the light intensity emitted from the LED is modulated to transmit data using a modified On-Off-Keying

¹In a more general case, a total of M LEDs can be used at each TX to satisfy the illumination level where the power consumed by each TX increases linearly with M .




Figure 2: Operating modes

Figure 3: LED I-V curve

(OOK) modulation, where *variable swing levels* I_{sw} around the bias are adopted to represent the symbol HIGH and LOW. A large I_{sw} can increase the received signal strength at the RX, but also consumes more power and may generate more interference to other RXs. A current $I_h = I_b + I_{sw}/2$ is used to represent a HIGH symbol, while a current $I_l = I_b - I_{sw}/2$ represents a LOW symbol. This ensures that illumination mode and illumination+communication mode generate the same brightness. The current I_b depends on the desired illumination level. A graphical representation is given in Fig. 2.

To avoid flickering and ensure equal probability of the HIGH and LOW levels, Manchester encoding is employed where a transition $I_l \rightarrow I_h$ denotes a binary 0 and a transition $I_h \rightarrow I_l$ denotes a binary 1. As a result, the same brightness of the LED is realized in the two operation modes.

Let $\mathbf{P}_t = (P_t^1, P_t^2, \dots, P_t^N)^T$ denote the optical power of the TXs for communication, and P_r the received optical power at a RX, then from Eq. (1) we have

$$P_r = \mathbf{H} \cdot \mathbf{P}_t + n, \quad (3)$$

where $\mathbf{H} = (H_1, H_2, \dots, H_N)$ is the line-of-sight path loss that can be derived from Eq. (2).

Objective: Let SINR_i be the received signal-to-interference-plus-noise ratio at RX i , and $P_{C,tot}$ the total power consumption for communication of all the TXs averaged over time. $P_{C,tot}$ can be expressed as:

$$P_{C,tot} = \sum_{j=1}^N \bar{P}_C^j, \quad (4)$$


where \bar{P}_C^j is the power consumed by TX j for communication averaged over time. Let $\mathbf{I}_{sw} = (I_{sw}^1, I_{sw}^2, \dots, I_{sw}^N)^T$ denote the swing level of the TXs. Our objective in this work is to select the swings \mathbf{I}_{sw} that optimize the sum log throughput, ensuring user proportional fairness [11]:

$$\max_{\mathbf{I}_{sw}} \sum_{i=1}^M \log(B \log_2(\text{SINR}_i + 1)) \quad (5)$$

subject to

$$0 \leq \sum_{k=1}^M I_{sw}^{j,k} \leq I_{sw,max}, \quad j = 1, \dots, N. \quad (6)$$

$$P_{C,tot} = \sum_{j=1}^N r \left(\frac{\sum_{k=1}^M I_{sw}^{j,k}}{2} \right)^2 \leq \bar{P}_{C,tot} \quad (7)$$

Figure 4: Approximation error on power consumption vs. swing level, with $I_b = 450\text{mA}$ (CREE XT-E LED)

where M denotes the number of RXs, $I_{sw,max}$ the maximal swing of the LED and $\bar{P}_{C,tot}$ the total power budget for communication. The inequality imposes both a minimum bound on the swing to ensure non-negativity and a maximum bound to ensure to operate in the communication region (derived further in Eq. (9)).

3.4 Analysis

To solve the optimization function (5) presented above, we need to derive the expressions for $P_{C,tot}$ and SINR_i .

3.4.1 Extra power consumption for communication. VLC consumes extra power beyond that for pure illumination [30]. To quantify the extra power \bar{P}_C^i consumed by each LED transmitter i for communication, we use the following equation that models the power consumption of an LED as a function of the current I that flows from the anode to the cathode:

$$P_{led}(I) = k V_t \ln \left(\frac{I}{I_s} + 1 \right) I + R_s I^2, \quad (8)$$

where k is the diode ideality factor, V_t the thermal voltage, I_s the reverse bias saturation current and R_s the series resistance of the LED [26]. As the modulation is performed around the bias current I_b , the Taylor expansion² of Eq. (8) around I_b gives the following expression:

$$P_{led}(I) \approx \underbrace{P_{led}(I_b)}_{P_I} + \underbrace{\frac{P'_{led}(I_b)}{1!} (I - I_b) + \frac{P''_{led}(I_b)}{2!} (I - I_b)^2}_{P_C}, \quad (9)$$

where the first term $P_{led}(I_b)$ represents the power consumption for illumination P_I , and the second and third terms refer to the power consumption P_C for communication. Since we use Manchester coding, the probabilities of symbol HIGH and symbol LOW in the modulated data are equal. The average power consumed on communication over time \bar{P}_C is:

$$\bar{P}_C = \mathbb{E}[P_C(t)] = \frac{kV_t}{2I_b} + 2R_s \frac{E[(I(t) - I_b)^2]}{2!} = r \left(\frac{I_{sw}}{2} \right)^2, \quad (10)$$

with $r = \frac{kV_t}{2I_b} + R_s$ the LED's dynamic resistance at the bias working point I_b . A graphical representation is shown in Fig. 3. Note that when the LED is used purely for illumination, the swing $I_{sw} = 0$ and

²Note that we consider the Taylor expansion up to the second-order term for the approximation in Eq. (9), sufficient for small variations around I_b .




Figure 5: Illumination distribution




Figure 6: 100 random instances




Figure 7: An illustrated instance

$\bar{P}_C = 0$, confirming our derivation. From Eq. (10), we can obtain the total extra power consumption for communication in the channel:

$$P_{C,tot} = \sum_{i=1}^N \bar{P}_C^i = \sum_{i=1}^N r \left(\frac{I_{sw}^i}{2} \right)^2. \quad (11)$$

Notice that the Taylor expansion is valid for small excitations around I_b . However, since the LED I-V curve can be approximated by a second order function, the approximation is also valid for larger swing levels, as shown in Fig. 4. Using the maximum swing level of $I_{sw} = 900\text{mA}$ results in an error of only 0.45%, which is acceptable.

3.4.2 Signal-to-interference-plus-noise ratio. The SINR at RX i is computed as follows:

$$\text{SINR}_i = \frac{\left(R\eta r \sum_{j=1}^N H_{j,i} \left(\frac{I_{sw}^{j,i}}{2} \right)^2 \right)^2}{N_0 B + \left(R\eta r \sum_{\substack{k=1 \\ k \neq i}}^M \sum_{j=1}^N H_{j,i} \left(\frac{I_{sw}^{j,k}}{2} \right)^2 \right)^2}, \quad (12)$$

with R the responsivity of the photodiode, η the wall-plug efficiency of the LED (i.e., the energy conversion efficiency from electrical to optical power), N_0 the single-sided spectral power density and B the communication bandwidth.

Note that the bias current I_b is not included to calculate the SINR, since it does not contribute to data transmission.

Since we have derived the expressions for $P_{C,tot}$ and SINR_i , then we can solve the optimization function (5) with nonlinear programming tools, such as `fmincon` in Matlab. Note that inequality (7) depends on the bias current I_b necessary to guarantee the required illuminance of the environment. Setting the bias I_b at the center of the linear region of the LED allows us to use a larger $I_{sw,max}$ (cf. Fig. 3). The opposite holds for a smaller or larger value of I_b , as the Taylor approximation will only be valid in a smaller region.

4 INSIGHTS INTO THE SYSTEM DESIGN

In this section, we provide an in-depth study of the system model introduced in Sec. 3, and present the main insights that will drive us in the design of a practical D-MISO system.

Simulation setup. We consider an indoor area of $3\text{m} \times 3\text{m} \times 2.8\text{m}$ consisting of a grid of $N = 36$ TXs and $M = 4$ RXs. The TXs are aligned in a 6×6 array with an inter-node distance of 0.5 m. They

Table 1: System parameters for simulations

Parameter	Notation	Value
General		
Single-sided spectral power density	N_0	$7.02 \times 10^{-23} \text{ A}^2/\text{Hz}$
Communication bandwidth	B	1 MHz
LED		
Half power semi-angle	$\phi_{1/2}$	15°
Reverse bias saturation current	I_s	$1.44 \times 10^{-18} \text{ A}$
Ideality factor, series resistance	k, R_s	2.68, 0.19 Ω
Bias current, wall-plug efficiency	I_b, η	450 mA, 0.4
Maximum swing current	$I_{sw,max}$	900 mA
Wall-plug efficiency	η	0.40
Receiver		
Field of view, collection area	Ψ_c, A_{pd}	$90^\circ, 1.1 \text{ mm}^2$
Responsivity	R	0.40 A/W

are attached at a height of 2.8m from the ground, facing downwards. The 4 RXs are located at a table of height 0.8m, facing upwards to the ceiling. The LED data is based on the off-the-shelf CREE XT-E LED [3], which is also used in our experimental validation in Sec. 8. Table 1 lists additional relevant system parameters.

Illuminance distribution. Fig. 5 shows the spatial illuminance distribution. According to the ISO 8995-1 illuminance standard for indoor office premises [5], the average illuminance should not fall below 500 lux and the illuminance uniformity (the ratio between the minimum illuminance and the average illuminance) shall not be less than 70%. We define an area of interest of $2.2\text{m} \times 2.2\text{m}$ (i.e. excluding the boundary), positioned in the center of the room. In this area, the average illuminance is 564 lux and the uniformity equals 74%, satisfying the illumination requirements.

4.1 Throughput vs. power consumption

Without loss of generality, we assume that the 4 RXs are located around certain TXs in the XY-plane, as shown in Fig. 6. For each RX, we generate 100 random positions around these TXs. We gradually increase the power budget $\bar{P}_{C,tot}$ for communication, and derive the optimal swings that lead to the maximal sum log throughput, by solving the optimization problem in Eq. (5). The results of

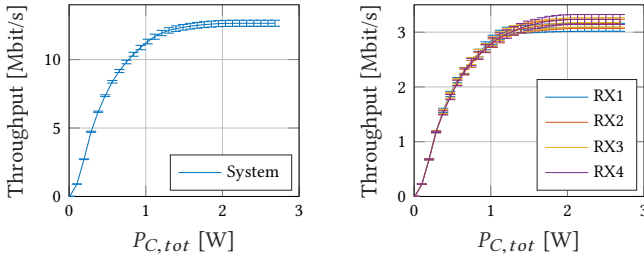


Figure 8: Average throughput over positions with 95% confidence interval versus communication power

the system throughput and each RX's throughput, both with 95% confidence interval, are depicted in Fig. 8.

We can observe that either the system or the individual's throughput increases in accordance with $\tilde{P}_{C,tot}$. Since we optimize the sum log throughput, the throughput of the RXs is balanced independent of $\tilde{P}_{C,tot}$, ensuring user fairness. When $\tilde{P}_{C,tot}$ is large, RX3 and RX4 achieve better performance than RX1 and RX2, due to the availability of more non-interfering TXs. However, we should note that the power efficiency on communication decreases when $\tilde{P}_{C,tot}$ exceeds a certain threshold. For instance, we can notice in Fig. 8 that the system throughput increases more slowly with the same extra power consumption when $\tilde{P}_{C,tot}$ exceeds 1.2 W.

4.2 Optimal swing levels

To show the results clearly, we consider a random instance of the RXs' positions used in the previous section. For this instance, the positions of the four RXs are shown in Fig. 7. The optimal swing levels leading to the maximized system throughput under different power budget $\tilde{P}_{C,tot}$ are shown in Fig. 9. Due the space limitation, we only show the results related to RX1 and RX2. We can observe that with the increase of $\tilde{P}_{C,tot}$, TX8 and TX10 are assigned first to RX1 and RX2, respectively, i.e. the swing level of a TX that poses the best channel increases first. When TX8 and TX10 reach the maximal allowed swing, the consumed communication power equals $P_{C,tx,max} = r \left(\frac{I_{sw,max}}{2} \right)^2 = 74.42$ mW, where r denotes the dynamic resistance. Increasing $\tilde{P}_{C,tot}$ further, more TXs start sequentially from zero-swing to operate at full-swing. For example to serve RX1, a group of TXs, forming the same CFM-MIMO beamspot, operate at full swing in the order of TX8→TX14→TX7→TX2→TX1→TX13.

Notice that the transition of a TX's optimal swing level from zero-swing to full-swing is fast. As shown in Fig. 9, with the increase of $\tilde{P}_{C,tot}$, the occurrence of the gray areas (representing that the TXs are operating at neither zero-swing nor full-swing) is negligible compared to the occurrence of the black (representing zero-swing) and white (representing full-swing) areas. This phenomenon applies to the 100 instances shown in Fig. 6.

For four representing TXs (TX3, TX5, TX10, and TX15), we depict the empirical CDF in Fig. 10 for five randomly selected instances³. TX10 has a very steep edge in the CDF at the maximum swing

³For better visualization, we only show the results for five instances.

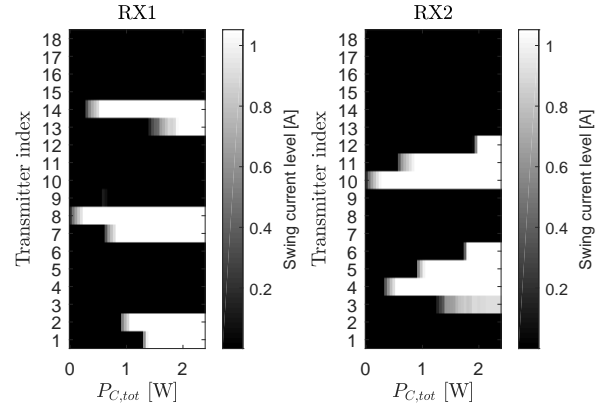


Figure 9: Swing levels vs. communication power

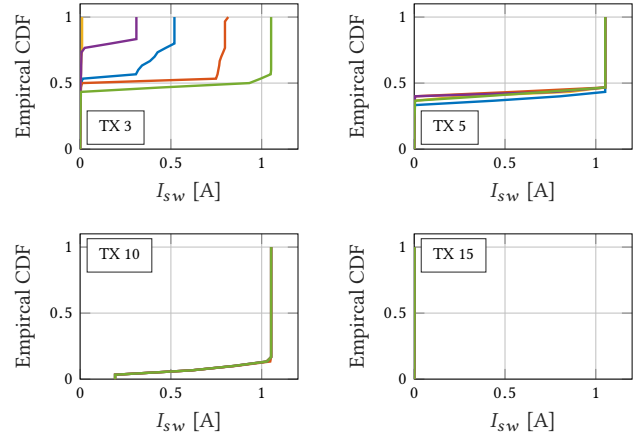


Figure 10: Empirical CDF vs. swing levels of RX2

level. This means that TX10 mostly operates at full-swing because it possess the best channel to RX2 among all the TXs. A similar trend is observed for TX5, although it is assigned later because it is farther from RX2 compared to TX10. Therefore, we see an offset in the CDF with respect to TX10.

An interesting observation in Fig. 10 is the empirical CDF for TX3. Instead of being a step function, the CDF of the optimal swing level of TX3 increases more smoothly and does not reach always the maximal swing, meaning that optimally it operates more often than other TXs at neither zero-swing nor full-swing. However, discretizing TX3's swing level to either zero or maximum results in a little loss with respect to the optimal swing level. The gain contributed by TX3 to the total system throughput is rather limited: if it is not used, the average loss in the total system throughput is only 0.48%. Finally, TX15 is not used at all because it would generate too much interference to the RXs.

Based on the above presented characteristics of the optimal swing levels, we have the following three insights:

Insight 1: Given a power budget for communication, for each RX, the optimal solution always assigns the available power sequentially to its preferred TXs. This means that the optimal solution

starts to assign the remaining communication power to the next preferred TX, only after the swing level of the RX's first preferred TX reaches the full-swing.

Insight 2: With the increase of the power budget, the transition of a TX's optimal swing level from zero-swing to full-swing is fast. That is, the probability of a TX operating at neither zero-swing nor full-swing are negligible compared to the probability of operating at either zero-swing or full-swing. Therefore, to achieve a near-optimal throughput in practice, it is sufficient for each TX to either use (i) $I_{sw}=0$ if operating in illumination mode, or (ii) the maximum swing $I_{sw,max}$ if operating in illumination+communication mode.

Insight 3: If the assignment of a TX to serve an intended RX generates much interference to other RXs, then that TX should be given lower priority in the sequence for power allocation. This means the TX will be assigned later, or not at all.

5 HEURISTIC ALGORITHM

Solving the optimization problem in Eq. (5) takes 165 seconds in Matlab, under the setup of 36 TXs and 4 RXs as presented in Sec. 4. This is not acceptable for a practical communication system. To reduce this complexity and facilitate a practical D-MISO system, in this section, we propose a ranking-based heuristic inspired by the insights gained in Sec. 4.

The key in our proposed heuristic is the custom-defined *Signal-to-Jamming Ratio (SJR)*. Its definition is inspired by Insight 3. Let the $N \times M$ matrix **SJR** denote the SJRs of all combinations of TXs and RXs:

$$\mathbf{SJR} = \begin{bmatrix} \text{SJR}_{1,1} & \cdots & \text{SJR}_{1,M} \\ \cdots & \text{SJR}_{n,m} & \cdots \\ \text{SJR}_{N,1} & \cdots & \text{SJR}_{N,M} \end{bmatrix}, \quad (13)$$

where $\text{SJR}_{i,j}$ of the TX i to the RX j is defined as:

$$\text{SJR}_{i,j} = \frac{H_{i,j}^\kappa}{\sum_{j'=1}^M H_{i,j'}}, \quad (14)$$

with κ a parameter that can be used to tune the weight of the desired channel to RX j with respect to the interference generated at other RXs. The higher the κ , the more weight is given to channel of TX i to RX j , and the less attention is paid to the interference generated by TX i to other RXs.

The proposed heuristic is outlined in Algorithm 1. It works as follows: after calculating all the SJRs, i.e. obtaining the **SJR** (line 1 to line 3), the heuristic searches sequentially for the TX having the maximum SJR (line 5). The tuple $\langle i^*, j^* \rangle$ that correspond to the maximum SJR is recorded and added to the ranking list *RankedTX* (line 6), which is a $1 \times N$ vector where its k th element RankedTX_k is a tuple $\langle i, j \rangle$, and $i \in \mathcal{N}, j \in \mathcal{M}$. Afterwards, the entire i th row of **SJR** is removed from the search space (line 7), since the i th TX is already assigned to a RX. This process is repeated until all TXs are added to ranking list *RankedTX*, i.e. when the search space of **SJR** is empty.

The proper value of the parameter κ depends on the system setup, such as the density of the TXs and the positions of the RXs. To verify the performance of our proposed heuristic, we set κ to different values and present the result of the ranking algorithm in Fig. 11. For a lower κ , i.e. $\kappa = 1.0$, the TXs give higher priority to

Algorithm 1: Proposed ranking-based heuristic

Input : $\kappa; H_{i,j}$
Output: *RankedTX*: a $1 \times N$ vector

```

1 for  $i \leftarrow 1$  to  $N$  do
2   for  $j \leftarrow 1$  to  $M$  do
3      $\text{SJR}_{i,j} \leftarrow \frac{H_{i,j}^\kappa}{\sum_{j'=1}^M H_{i,j'}}$ 
4 for  $k \leftarrow 1$  to  $N$  do
5    $(i^*, j^*) \leftarrow \arg \max_{i,j} \text{SJR}$ 
6    $\text{RankedTX}[k] \leftarrow (i^*, j^*)$ 
7    $\mathbf{SJR} \leftarrow \mathbf{SJR} \setminus \text{SJR}_{i^*,:}$ 

```

avoid generating interference to other unintended RXs, therefore, resulting a low system throughput when the power budget is low. For a larger value of κ , e.g. $\kappa = 1.5$, the TXs prefer to serve the intended RXs, which generates more interference when the budget is large, also resulting in a lower throughput. $\kappa = 1.2$ and $\kappa = 1.3$ show a good balance between the channel quality and generated interference. Compared to the optimal system throughput gained by solving the optimizing problem, the average throughputs with $\kappa = 1.0, \kappa = 1.2, \kappa = 1.3, \kappa = 1.5$ are decreased by 40.3%, 2.4%, 1.8%, 2.6%, respectively. The histogram of average loss in system throughput with different κ for the 100 random instances (cf. Fig. 6) is further shown in Fig. 11. So in this setup, $\kappa = 1.3$ is the best choice.

Note that our proposed ranking-based heuristic only takes 0.07 seconds, i.e. reducing the complexity by 99.96% compared to the optimal solution presented before, at the cost of scarifying only 1.8% of the system throughput. Besides, our heuristic can be easily implemented in a micro controller.

6 SYNCHRONIZATION

For each CFM-MIMO beamspot in DenseVLC, there is a varying number of TXs that are dynamically synchronized to send the same information signal to an RX. Thus, fine-grained, flexible and scalable synchronization among the TXs is very important. In this section, we first state the problems of using existing protocols in DenseVLC, then propose a novel and practical synchronization method based on non-line-of-sight (NLOS) VLC.

6.1 Synchronization with NTP/PTP

We first test two synchronization protocols together in DenseVLC: Network Time Protocol (NTP) [6] and Precision Time Protocol (PTP) [4]. NTP is used to coarsely synchronize the controller's time to the time of an external server, while PTP is in charge for a finer time synchronization among VLC TXs. In addition, in this test, the frame structure of DenseVLC (cf. Table 3) is modified such that each TX will start transmission at the absolute time specified in μs in the Sync field, appended to the frame. Knowing this *absolute time* to transmit, each TX waits until the specified absolute time after which they transmit the data simultaneously.

We implement the above approach and measure the synchronization delay between two received signals transmitted simultaneously

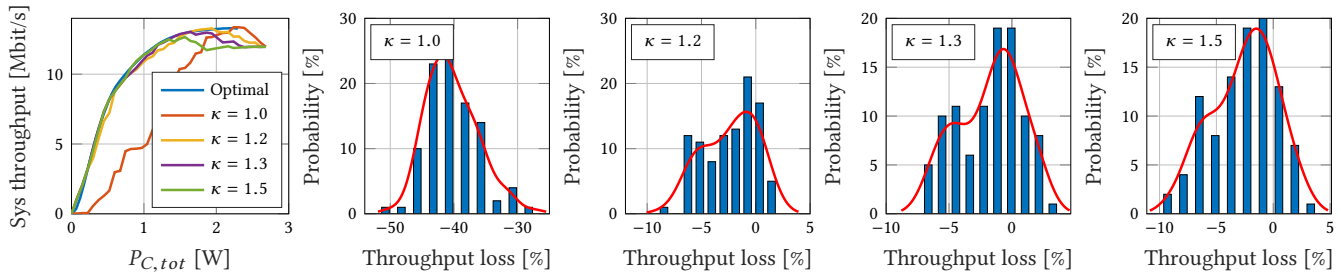


Figure 11: Heuristic verification: left) system throughput obtained with heuristic, for the instance shown in Fig. 7; right) average loss in system throughput, for the 100 random instances shown in Fig. 6




Figure 12: Synch delay **Figure 13: Synch with NLOS VLC**

by two TXs. The time difference of each two ‘synchronized’ symbols from the two TXs is computed, and the median delay over an entire frame is derived. This process is repeated 10 times (for 10 different frames) and we average out the 10 median delays as the final synchronization delay. We measure the synchronization delay at several different symbol rates, for which the results are shown in Fig. 12. We observe that the NTP/PTP-based synchronization method significantly improves the delay, by at least a factor of two between two VLC TXs. If the maximum acceptable overlap between ‘synchronized’ symbols is 10% of the symbol width, then the maximum symbol rate at which we can transmit with two synchronized TXs is 14.28 Ksymbols/s.

Therefore, *this approach has a fundamental limitation in symbol rate*. The TXs cannot be synchronized with a higher accuracy than the one showed above, because it relies on external libraries running on top of an operating system. This approach could work for low transmission rates, but is unreliable for faster communication. Moreover, it requires CPU time and synchronization messages through Ethernet, which can further degrade the synchronization accuracy when the network workload is higher.

6.2 Synchronization with NLOS VLC

In DenseVLC, we propose a novel synchronization method by exploiting NLOS VLC. For each RX, the controller appoints a leading TX from those that will jointly serve the RX. The leading TX transmits the pilot signal to synchronize the rest of the TXs. Then they jointly send synchronized signal to the RX. This method can reduce synchronization delay and does not rely on external libraries running on top of an operating system. It uses relative time instead of absolute time, and therefore, does not need to connect to an external time server. The proposed method works as follows:

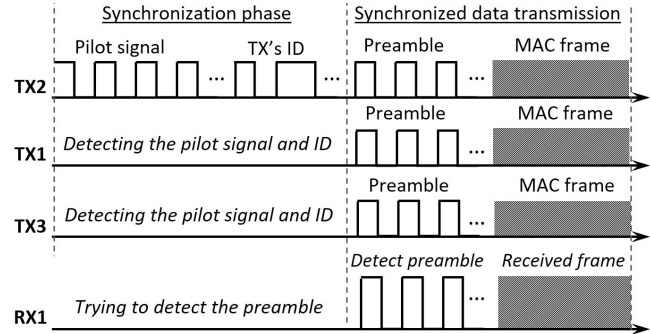


Figure 14: The proposed method that exploits NLOS VLC for wireless synchronization

- For each RX, the controller sends data to the desired TXs (resulting from our heuristic) and appoints one of them as the leading TX. The leading TX, e.g. TX2 as illustrated in Fig. 14, first transmits the pilot signal for synchronization and its own ID, then transmits the real frame (Preamble + MAC frame). The symbol rate at TX2 is denoted as f_{TX} .
- The non-leading TXs, e.g. TX1 and TX3 in Fig. 14, listen for the synchronization signal with a sampling rate of f_{RX} that is much higher than f_{TX} . The higher the f_{RX} compared to f_{TX} , the better synchronization we can achieve.
- After detecting the pilot signal and the leading TX’s ID, RXs compare the ID to their desired one. If matched, they start sending synchronously the preamble and MAC frame after a short pre-defined guard period.

The proposed synchronization method is scalable in terms of the number of TXs and RXs because the synchronization is performed only by listening to the pilot signal from nearby leading TXs. The synchronization method is also flexible since we rely on wireless visible light communications and do not need to wire each two TXs physically to achieve good synchronization.

7 SYSTEM IMPLEMENTATION

In this section we present the implementation of DenseVLC, including both the hardware and software implementations.




Figure 15: Design of the TX front-end

7.1 Hardware

We use the Beaglebone Black (BBB) as embedded computer in our design. The controller has a dedicated BBB to orchestrate the CFM-MIMO beamspots and run the VLC MAC. Each TX has a front-end for transmission, and a receiving front-end for synchronization. The VLC PHY of four TXs is managed by 1 BBB, so 9 BBBs are used in total for the 36 TXs. For each RX, its front-end is mounted on a dedicated BBB Wireless. Next, we present the two front-end designs.

TX front-end: We use the high performance LED CREE XT-E, covered by a lens to limit the field of view of the LED. We design a circuit to emit light with three different intensity levels (symbol LOW, illumination, symbol HIGH), instead of two levels (light and no light) as done in typical low-end VLC systems [32]. In our design, the LED emits no light when transmitting the symbol LOW. For illumination and symbol HIGH, we use two parallel branches, containing a power transistor and a resistor in series, to drive the LED. The two resistor values are tuned such that the average luminous flux from the LED does not change when going from illumination mode to 50% duty cycled communication mode and vice versa. The average measured electrical power consumption is 2.51W for illumination and 3.04W for 50% duty cycled communication. The TX front-end supports a transmission rate up to 2MHz. The developed PCB board and block diagram of the TX front-end are presented in Fig. 15.

RX front-end: We design a three stage analog RX front-end to amplify the signal with minimal noise. In the first stage, the photodiode current is amplified to a voltage by a low noise Transimpedance amplifier (TIA). The second stage contains an AC coupled amplifier, which filters out low frequency ambient light and amplifies the signal further, and thus enables the detection of a very weak signal, e.g., a reflected pilot signal. In the third stage, the high frequency noise components are removed by a 7th order passive low pass Butterworth filter, to avoid aliasing in the subsequent analog-to-digital converter (ADC). The ADC digitizes the signal at 1 Msamples/s and sends it via SPI to the Programmable Real-Time Unit (PRU) of the BBB Wireless. The developed BBB cape and the block diagram of the RX front-end are presented in Fig. 16. Note that this front-end is used for both the receiver on the ground and the transmitter on the ceiling, but with different software; the former to receive data packets, and the latter to receive synchronization pilot signals.

7.2 Software

The DenseVLC software is developed partly based on the OpenVLC platform [10]. The source codes of DenseVLC have been made public

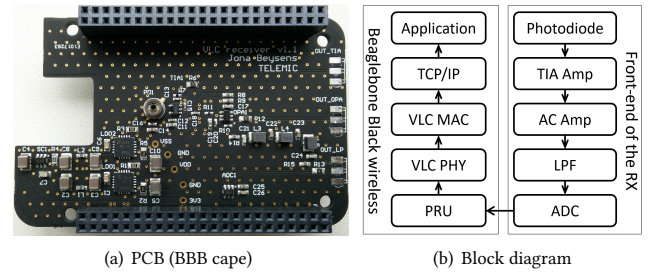


Figure 16: Design of the RX front-end

Table 2: Hardware components

TX		RX	
LED	CREE XT-E	Photodiode	S5971
Lens	TINA FA10645	TIA / AC Amp	OPA659 / OPA355
Power transistor	NTR4501	ADC	ADS7883

and can be found on https://github.com/jonabeyens/CoNEXT18_DenseVLC.⁴

Data communication from controller to TXs. After encapsulating the data with an Ethernet header, the frames are sent from the controller via Ethernet multicast to the TXs. A high level illustration of the frame structure is shown in Table 3. This header is decapsulated once it reaches the VLC TX. The start of the VLC PHY header contains an 8 bytes field specifying the IDs of the TXs that should transmit the data. Each TX checks this field and acts upon it accordingly. The RX receives the data and if the frame is decoded successfully, sends a MAC acknowledgement frame back to the controller using WiFi [17]. We have taken advantage of the fact that BBB wireless has WiFi connectivity already integrated. As such, the system has full mobility. Note that uplink packets are usually smaller in quantity and size compared to downlink packets. Therefore, the WiFi link is not easily congested.

Channel measurements. The controller selects the optimal operation mode of the VLC TXs based on periodic channel measurements. In order to quantify the channel quality, each TX sends sequentially a predefined bit stream to the RX. The RX measures received signal strength and sends this information to the controller. To minimize the resulting signaling overhead, this response is fit in a frame with minimal length and sent when the channel is idle. To estimate the SNR from the experimental data, the M2M4 estimator is used since (1) it shows good performance for AWGN channels and (2) it can use symbols of the frame after the ADC, without the need for estimating the channel before communication [23]. As the noise can differ during the reception (depending on the circuitry of the RX and the photodiode's negative bias), the latter property is useful to better characterize the SNR.

Decision logic. Once the controller has received the channel updates for all the links from the TXs to the RXs, it selects $I_{SW} =$

⁴This repository is also made available at https://github.com/openvlc/CoNEXT18_DenseVLC.

Table 3: Frame structure (Controller to VLC TXs; B: bytes)

ETH PHY + MAC header	TX ID	Pilot sig.	Preamble	SFD	Length	Dst	Src	Protocol	Payload	Reed-Solomon
	8B	32 symbols	32 symbols	1B	2B	2B	2B	2B	xB	$\lceil x/200 \rceil$ 16B

$\{0, I_{sw, max}\}$ for each VLC TX based on our ranking algorithm presented in Sec. 5.

8 PERFORMANCE EVALUATION

In this section we evaluate the performance of DenseVLC.

Experimental setup. We evaluate the performance of DenseVLC in a system of 36 VLC TXs and 4 RXs. The TXs are deployed 6×6 within an area of $3\text{m} \times 3\text{m}$, with 0.5 m inter-TX distance, and a height of 2 m from the floor. The 4 RXs are placed on the floor, controlled by 4 OpenBuilds ACRO System [2] and can be moved to any position within the $3\text{m} \times 3\text{m}$ area. The experimental setup is depicted in Fig. 17.

Illuminance distribution. DenseVLC provides an average illuminance of 530 lux and an uniformity of 81%. The measurements were performed with the HS1010 lux meter.

8.1 Synchronization evaluation

The synchronization between TXs is one of the key enablers in DenseVLC. To evaluate the proposed method that exploits NLOS VLC for synchronization, we first randomly choose two neighboring TXs, TX2 and TX3. TX2 is appointed as the leading TX to send the pilot signal for synchronization. The symbol rate f_{tx} at TX2 and the sampling rate f_{rx} at TX3 are set to 100 Ksymbols/s and 1 Msamples/s, respectively. We connect the anodes of the LEDs at TX2 and TX3 to an oscilloscope (RIGOL MSO1104) to capture the transmitted signals. The delay between the corresponding symbol edges of the two signals are measured. As in Sec. 6.1, we calculate the median of the synchronization delay and compare it to those from the method based on NTP/PTP and to the one without synchronization. The results are shown in Table 4. We can see that the median synchronization delay of our method with NLOS VLC is only $0.575 \mu\text{s}$, improving the synchronization granularity by nearly an order of magnitude compared to the one using NTP/PTP. Note

**Figure 17: Experimental setup****Table 4: Evaluation of the proposed synchronization**

	No Synchronization	NTP/PTP	NLOS VLC
Median error	$10.040 \mu\text{s}$	$4.565 \mu\text{s}$	$0.575 \mu\text{s}$

Table 5: Experimental result using iperf

Scenario	Throughput [Kbit/s]	PER [%]
2 TXs	33.9	0.19
4 TXs (no sync)	0	100
4 TXs (with our sync)	33.8	0.55

that with advanced devices supporting a higher sampling rate of f_{rx} , the synchronization granularity supported by our NLOS VLC based method can be further improved.

To test the synchronization performance, we perform iperf measurements for 100 seconds under three different scenarios. For all scenarios, there is one RX, located in the center of TX2, TX3, TX8 and TX9. The results are shown in Table 5. In the first scenario, only TX2 and TX8 serve the RX. Since TX2 and TX8 are managed by the same BBB, no synchronization is required. The packet error rate (PER) is low, due to the strong signal strength and low noise at the RX. The achieved throughput is lower than the used symbol rate of 100 Ksymbols/s, due to Manchester encoding, PHY and MAC layer overhead and Reed-Solomon error correcting. In the second scenario, TX3 and TX9, managed by another BBB, also serve the RX. However, no synchronization is enabled. No packets are received, due to improper alignment of the frames in time. In the last scenario, synchronization is added. Very low packet loss is observed again, showing that our NLOS VLC based synchronization works.

8.2 Heuristic evaluation

To evaluate our proposed ranking algorithm in DenseVLC, we carry out experiments under three representing scenarios, for which the RX positions are listed in Table 6:

- *Scenario 1:* interference-free; no dominating TX.
- *Scenario 2:* with interference; no dominating TX.
- *Scenario 3:* with interference; with dominating TX.

First, we perform experimental channel measurements from the 36 TXs to the 4 RXs. Afterwards, the path loss is computed as the received swing level at the RX and reported to the controller. Using the path loss data, the controller runs the ranking-based heuristic as presented in Algorithm 1 for different values of κ . We assign the TXs from the ranked list one by one (i.e. increasing the allowed power budget on communication step by step) to the corresponding RXs, and calculate the SINR based on Eq. (12) with the experimental data. The system throughput is obtained based on Eq. 5.




Figure 18: Experimental results for Scenario 1

Scenario 1: In this scenario, the inter-RX distance is 2 m. The result of the ranking algorithm is shown in Fig. 18. When we gradually increase the $P_{C,tot}$ by assigning more and more TXs, we observe that assigning a TX to one RX results in no throughput drop to the other RXs, which means the interference is limited. The system throughput for the different κ values is similar, although $\kappa = 1.0$ performs a little bit worse because with a smaller κ , a little bit of interference could prevent using more TXs for a RX.

Scenario 2: The RX positions are now identical to the simulation setup presented in Fig. 7. We observe that the experimental results show a similar trend for the system throughput as the simulations, presented in Fig. 11. As in the simulations, $\kappa = 1.3$ also shows good performance here. Further, we can observe from Fig. 19 that in the beginning the throughput of all the RXs increases at the same rate, but afterwards RX1 achieves a lower throughput than the other RXs. This is because RX1 is positioned closer to the TXs that generate most interference compared to the other RXs. Next, assigning a TX to one RX has more influence on the throughput of the other RXs than in Scenario 1, especially at higher $P_{C,tot}$. Finally, $\kappa = 1.0$ pays too much attention to interference at low $P_{C,tot}$, resulting in a lower system throughput.

Scenario 3: The inter-RX distance in this scenario is 1 m, and every RX is located exactly under 1 TX. From Fig. 20, we observe similar results as in Scenario 2. However, RX1 achieves now comparable throughput to the other RXs. Further, due to slightly more interference as in Scenario 2, the system throughput drops when assigning many TXs.

8.3 Power efficiency

In order to benchmark the performance of our method, we compare DenseVLC with following two techniques:

- *Nearest-TX communicating (SISO):* only the nearest TX to the RX is used for communication, resulting in 4 assigned TXs in total. The others are used solely for illumination.

Table 6: RX positions in the experiments (in meter)

Scenario	RX1	RX2	RX3	RX4
1	(0.50,0.50)	(2.50,0.50)	(0.50,2.50)	(2.50,2.50)
2	(0.92,0.92)	(1.65,0.65)	(0.72,1.93)	(1.99,1.69)
3	(0.75,0.75)	(1.75,0.75)	(0.75,1.75)	(1.75,1.75)

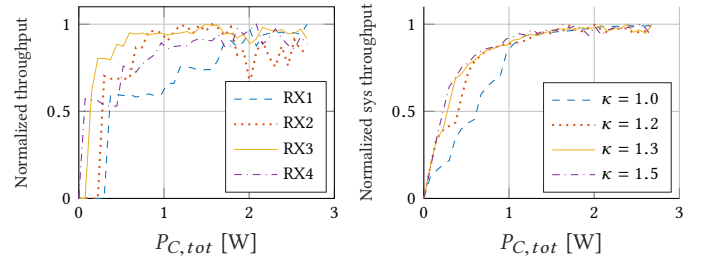


Figure 19: Experimental results for Scenario 2

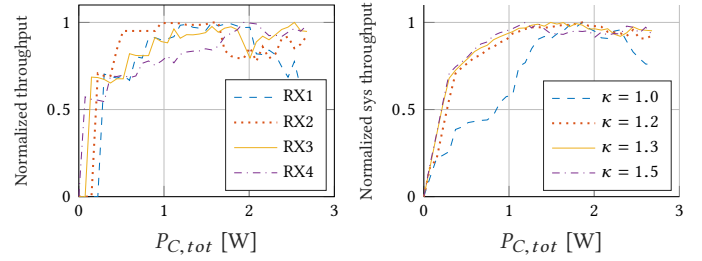


Figure 20: Experimental results for Scenario 3

- *All-TXs communicating (D-MISO):* all TXs are used for communication, independent of the position of the receivers. For this setup, this means that each RX is assigned 9 surrounding TXs.

Fig. 21 depicts the system throughput in Scenario 2 for DenseVLC with $\kappa = 1.3$, SISO and D-MISO. The markers denote the operating point of SISO and D-MISO, and a horizontal line is added to facilitate the comparison. We observe that the operating point of SISO crosses with DenseVLC, meaning that it achieves the same power efficiency, i.e. throughput versus power. The power consumption is 298mW, for a normalized throughput of 0.63. However, as opposed to SISO, DenseVLC can provide more throughput by adding extra TXs. By adding more TXs, the gap between DenseVLC and D-MISO reduces. At $P_{C,tot} = 1.19W$, DenseVLC achieves the same normalized throughput of 0.94 as D-MISO. However, D-MISO would require power consumption of 2.68W to achieve the same result. Therefore, DenseVLC can improve the power efficiency by 2.3 times. The gain in throughput with respect to SISO for this operating point equals 45%. Due to space limitations, we only show the result for Scenario 3, but the conclusion is also valid for the other scenarios.

9 LIMITATIONS AND DISCUSSIONS

We discuss current limitations of DenseVLC in this section.

Personalized and adaptive κ . In the heuristic presented in Sec. 5, we use the same κ for all the TXs when calculating the signal-to-jamming ratio. In a real cell-free massive MIMO system, TXs will cause different interference to unintended RXs, depending on RXs' positions and the surroundings. Therefore, properly personalized and adaptive κ s can boost the system performance towards the

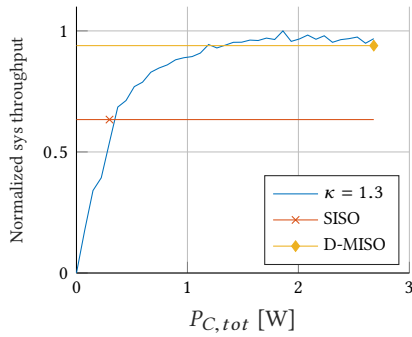


Figure 21: DenseVLC vs. SISO and D-MISO

optimal result obtained from solving the optimizing problem in Eq. (7).

Advanced hardware. One of the limitations in DenseVLC is the low-cost PRU/BBB used for sampling. With advanced dedicated hardware such as FPGA, we can perform the sampling much faster and as such achieve better synchronization granularity, as well as to exploit advanced modulation schemes such as OFDM in VLC.

Blockage. In traditional VLC, blockage can degrade the system performance greatly because it can break visible light communication links. In cell-free massive MIMO systems with VLC, however, blockage could bring benefit to the system since it can reduce the interference from other TXs. We will investigate this in DenseVLC in the near future.

TX and RX density. Ideally, a user is served by a transmitter which is as close as possible, because then the received desired signal is the highest, and the interference the lowest. The lower the TX density, the less degrees of freedom we have to serve the users. This results in a both a lower system throughput and user fairness. Next to this, also the impact of the number of users on the system performance. We will evaluate the impact of TX and RX density in the future work.

RX orientation. Although receiver orientation has an impact on the receiver's SINR, this is not studied here since it is not the focus of this work. Both the optimization problem and the heuristic are not limited to facing up receivers, and work for all receiver orientation.

NLOS synchronization. The proposed NLOS synchronization approach relies on the reflected pilot signals from the floor. From the preliminary study on the performance for different floor materials, we observe that the pilot signal can also be detected with less reflective floor materials. Further, the evaluation is done without human motion. However, we notice that even when a person is walking by, the pilot signals are still received and as such this does not harm the synchronization significantly. An extensive evaluation has been planned for future work.

10 RELATED WORK

We summarize the most relevant works in this section.

MIMO/(MU-)MISO in VLC. There is many research done on general MIMO schemes in VLC networks. Here, we only cover those that are most related to our work. In [34], the authors consider a MIMO system where each TX has multiple independent

LEDs. They design an algorithm that can select the best subset of the LEDs of each TX based on the channel qualities, to optimize the system capacity, while satisfying the required uniform illumination. In [37], the authors design precoders to improve SNR performance of a non-distributed MIMO system with constraint on the LED-dependent swing. They show that the larger the swing, the lower the achieved BER. The authors in [25] target at maximizing the SNR with the constraint on power consumption. The problem is formulated as a non-linear optimization and solved numerically. However, it cannot provide stable illumination. In [13], they introduce a cell-free VLC system that can achieve proportional fairness with three transmitters. Compared to the above work, we study a scenario with massive VLC TXs and multiple RXs, allowing us to build a real CFM-MIMO system. We design the system architecture, optimize throughput, design and build a testbed to evaluate the performance. We also propose a novel method to synchronize multiple distributed TXs with NLOS VLC. None of the above work consider the practical synchronization, and none of them has a complete system design and experimental evaluation.

Power-efficient VLC. VLC consumes extra power beyond illumination. The authors in [30] show that this extra power consumption depends on the LED, the illumination pattern, and the driver circuit. To have a power-efficient VLC system, these aspects can be exploited. For example, [9] exploits the illumination patterns to achieve an energy efficient system while the authors in [12] design an energy-efficient LED driver for VLC. In this work, our analysis on the extra power consumption for communication is inspired by [30] and our analytical results and findings guide us to design the system components of DenseVLC and build a real system.

11 CONCLUSION

Motivated by the dense-luminaries infrastructure and the concept of Cell-Free Massive MIMO, we proposed DenseVLC, a novel massive VLC system enabled by densely distributed LEDs and visible light communications. DenseVLC adapts the beamsports based on system dynamics to serve multiple users simultaneously. We designed a heuristic to reduce the complexity on adaptation and proposed to synchronize distributed TXs with NLOS VLC. We implemented DenseVLC with off-the-shelf devices and evaluated its performance in extensive experiments. Our results demonstrated the feasibility of building a practical CFM-MIMO system with VLC. Going forward, we envision that DenseVLC can inspire follow-up system research towards this promising direction.

ACKNOWLEDGMENT

This work is supported in part by the Research Foundation Flanders (FWO) SB PhD fellowship under grant number 1S59918N (Jona Beysens), the FWO postdoctoral fellowship under grant number 12Y0919N (Qing Wang), and the "La Caixa international PhD program" fellowship (Ander Galisteo). Our sincere thanks to the reviewers and our shepherd Yunfei Ma for all of their great comments and constructive feedbacks.

REFERENCES

- [1] http://www.earth-policy.org/data_highlights/2011/highlights15.
- [2] ACRO. <https://openbuilds.com/builds/openbuilds-acro-system.5416/>.
- [3] CREE XT-E LED. <https://goo.gl/ocs7bP>.

- [4] IEEE Std 1588-2008. *IEEE Standard for a Precision Clock Synchronization Protocol for Networked Measurement and Control Systems*.
- [5] ISO 8995-1:2002 - Lighting of work places – Part 1: Indoor.
- [6] RFC 5905. *Network Time Protocol Version 4: Protocol and Algorithms Specification*.
- [7] AZIZAN, L., AND ET. AL. Optimization of snr for wireless light-emitting diode communication in modern lighting layouts. *Optical Engineering* (2014).
- [8] DIMITROV, S., AND HAAS, H. *Principles of LED Light Communications: Towards Networked Li-Fi*. Cambridge University Press, 2015.
- [9] DIN, I., AND KIM, H. Energy-efficient brightness control and data transmission for visible light communication. *IEEE Photonics Technology Letters* (2014).
- [10] GALISTEO, A., JUARA, D., WANG, Q., AND GIUSTINIANO, D. Openvlc1.2: Achieving higher throughput in low-end visible light communication networks. In *Wireless On-demand Network Systems and Services (WONS)* (2018).
- [11] GIUSTINIANO, D., GOMA, E., LOPEZ TOLEDO, A., DANGERFIELD, I., MORILLO, J., AND RODRIGUEZ, P. Fair wlan backhaul aggregation. In *MobiCom* (2010).
- [12] JALAJAKUMARI, A., AND ET. AL. An energy efficient high-speed digital led driver for visible light communications. In *ICC* (2015).
- [13] JIANG, R., WANG, Q., HAAS, H., AND WANG, Z. Joint user association and power allocation for cell-free visible light communication networks. *IEEE Journal on Selected Areas in Communications* (2018).
- [14] KOMINE, T., AND NAKAGAWA, M. Fundamental analysis for visible-light communication system using led lights. *IEEE transactions on Consumer Electronics* (2004).
- [15] LOMBARDI, M. A., NELSON, L. M., NOVICK, A. N., AND ZHANG, V. S. Time and frequency measurements using the global positioning system. *Cal Lab: International Journal of Metrology* (2001).
- [16] MORENO, I., AVENDA, M., AND TZONCHEV, R. Designing light-emitting diode arrays for uniform near-field irradiance. *Appl. Opt.* (2006).
- [17] NARIBOLE, S., AND ET. AL. LiRa: A WLAN Architecture for Visible Light Communication with a Wi-Fi Uplink. In *IEEE SECON* (2017).
- [18] NAYEBI, E., ASHIKHMIN, A., MARZETTA, T., AND YANG, H. Cell-Free Massive MIMO systems. In *Asilomar Conference on Signals, Systems and Computers* (2015).
- [19] NAYEBI, E., ASHIKHMIN, A., MARZETTA, T., YANG, H., AND RAO, B. Precoding and Power Optimization in Cell-Free Massive MIMO Systems. In *IEEE Transactions on Wireless Communications* (2017).
- [20] NGO, H., ASHIKHMIN, A., YANG, H., LARSSON, E., AND MARZETTA, T. Cell-Free Massive MIMO Versus Small Cells. In *IEEE Transactions on Wireless Communications* (2017).
- [21] NGO, H., TRAN, L., DUONG, T., MATTHAIYOU, M., AND LARSSON, E. On the total energy efficiency of cell-free massive mimo. In *IEEE Transactions on Green Communications and Networking*.
- [22] NGUYEN, L., DUONG, T., NGO, H., AND TOURKI, K. Energy Efficiency in Cell-Free Massive MIMO with Zero-Forcing Precoding Design. In *IEEE Communications Letters* (2017).
- [23] PAULUZZI, D., AND BEAULIEU, N. A comparison of SNR estimation techniques for the AWGN channel. *IEEE Trans. on Communications* (2000).
- [24] RAMANE, D., AND SHALIGRAM, A. Optimization of multi-element led source for uniform illumination of plane surface. *Optics express* (2011).
- [25] SAENGUDOMLERT, P. Transmit beamforming for line-of-sight mimo vlc with im/dd under illumination constraints. In *IEEE ECTI-CON* (2015).
- [26] SCHUBERT, E. *Light-Emitting Diodes*. Cambridge University Press, 2006.
- [27] SEWAIWAR, A., AND ET. AL. Smart led allocation scheme for efficient multiuser visible light communication networks. *Optics express* (2015).
- [28] SU, Z., XUE, D., AND JI, Z. Designing led array for uniform illumination distribution by simulated annealing algorithm. *Optics express* (2012).
- [29] SUN, Z., YU, H., LI, W., TIAN, Z., AND ZHU, Y. Power-Efficient Linear Precoding for MU-MISO VLC Systems With Channel Uncertainty. In *IEEE Photonics Technology Letters* (2018).
- [30] TSIATMAS, ET AL, A. Joint illumination and visible-Light Communication systems: Data rates and extra power consumption. *IEEE ICC* (2015).
- [31] VARMA, G., AND ET. AL. Power allocation for uniform illumination with stochastic led arrays. *Optics Express* (2017).
- [32] WANG, Q., GIUSTINIANO, D., AND PUCCINELLI, D. Openvlc: Software-defined visible light communication networks. In *ACM VLCS* (2014).
- [33] WANG, Z., AND ET. AL. Performance of a novel LED lamp arrangement to reduce SNR fluctuation for multi-user VLC systems. *Optics express* (2012).
- [34] WANG, Z., GUO, C., YANG, Y., AND LI, Q. Antenna selection based dimming scheme for indoor MIMO visible light communication systems utilizing multiple lamps. In *IEEE PIMRC* (2016).
- [35] YAMADA, M., AND CHWASTYK, D. Adoption of light-emitting diodes in common lighting applications. Tech. rep., 2015.
- [36] YANG, H. A. Uniform illumination rendering using an array of leds: a signal processing perspective. *IEEE trans. on signal processing* (2009).
- [37] YING, K., AND ET. AL. MIMO transceiver design in dynamic-range-limited vlc systems. *IEEE Photonics Technology Letters* (2016).
- [38] ZHANG, J., ZHANG, X., AND WU, G. Dancing with light: Predictive in-frame rate selection for visible light networks. In *INFOCOM* (2015).
- [39] ZHANG, Y., YU, H., AND ZHANG, J. Block precoding for peak-limited miso broadcast vlc: Constellation-optimal structure and addition-unique designs. In *IEEE Journal on Selected Areas in Communications* (2018).