Dancing with Light: Predictive In-frame Rate Selection for Visible Light Networks

Jialiang Zhang
University of Wisconsin-Madison & UESTC
jialiang.zhang@wisc.edu

Xinyu Zhang
University of Wisconsin-Madison
xyzhang@ece.wisc.edu

Gang Wu
UESTC, China
wugang99@uestc.edu.cn

***Abstract*—Visible Light Communications (VLC) is emerging as an appealing technology to complement WiFi in indoor environments. Yet maintaining VLC performance under link dynamics remains a challenging problem. In this paper, we build a VLC software-radio testbed and examine VLC channel dynamics through comprehensive measurement. We find minor device movement or orientation change can cause the VLC link SNR to vary by tens of dB even within one packet duration, which renders existing WiFi rate adaptation protocols ineffective. We thus propose a new mechanism, DLit, that leverages two unique properties of VLC links (predictability and full-duplex) to realize fine-grained, in-frame rate adaptation. Our prototype implementation and experiments demonstrate that DLit achieves near-optimal performance for mobile VLC usage cases, and outperforms conventional packet-level adaptation schemes by multiple folds.**

I. Introduction

The concept of Visible Light Communications (VLC) originates from optical wireless technology which has already been adopted in mission-critical communication scenarios, *e.g.*, deep-space and military cases. Yet recent years witnessed a renewed interest in consumer-grade VLC, primarily because of the looming spectrum crunch in the legacy RF domain, and the increasing installation of LED lights, which are expected to phase out the power-hungry incandescent/fluorescent lamps by 2018 [1], [2]. A VLC system delivers digital signals by controlling the ON/OFF repetition of LED transmitters at a much faster speed ($>$ 100Hz) than the persistence of human eye. A photodiode (PD) receiver can sense the light-intensity variation and subsequently recover the digital information.

VLC possesses a number of appealing advantages over existing communications technologies. It has a well confined communication coverage and resists eavesdropping, and introduces no interference to RF devices. Thus, it holds great potential to alleviate the spectrum crunch in current wireless access networks. It can enable a wide range of indoor short-range communications, *e.g.*, between ceiling LEDs and smartphones/wearables, apparatus in hospitals (where RF is prohibited) [3]. The IEEE 802.15.7 has standardized VLC MAC/PHY protocols in order to spur the development of devices to support such real-world applications. Substantial recent research effort has been devoted to developing devices and modulation schemes to boost the VLC bit-rate [1], [2].

Despite the encouraging opportunities from a communications perspective, grand challenges remain in networking VLC devices. In particular, *how to ensure robust connectivity between the VLC transmitter and receiver, in the presence of link dynamics (e.g., device misalignment, mobility)?* This involves a myriad of networking problems, such as channel access, node discovery and handoff between LED cells.

New challenges in rate adaptation. This paper focuses on a specific aspect of VLC networking, *i.e.*, link rate adaptation. We consider indoor VLC links between an LED fixture (*e.g.*, ceiling-mounted lamp) and a PD equipped on mobile handset or moving object (*e.g.*, wheel-chair or robot cleaner). Rate adaptation has been extensively investigated in 802.11-based WiFi networks. However, through testbed measurement, we find that *a VLC link exhibits highly dynamic spatial/temporal variations that can no longer be tackled by conventional rate adaptation protocols*.

A wireless link may harness signal diversity from multipath reflections — abundant strong signals can be received even when a device is slightly rotated or moved. In contrast, VLC heavily relies on a single line-of-sight (LOS) path and the received light intensity is highly sensitive to the orientation between transceivers. Our experiments reveal that for a VLC receiver, an angular change of 30 degrees, or displacement of 8 centimeters, can cause around 10 dB of SNR variation. When it is held by a walking person, the SNR varies by up to 10 dB within a few milliseconds. Given the VLC bit-rate as defined in IEEE 802.15.7, a packet may last several hundred milliseconds – much longer than the channel coherence time. Therefore, conventional rate adaptation protocols, which generally rely on packet-level statistics from historical receptions, lose their efficacy in VLC networks.

DLit's principles. To meet the new challenges, we propose a new rate adaptation mechanism called *DLit* that leverages two unique properties in VLC: channel predictability and full-duplex. *First*, our empirical study indicates that the first and second-order statistics of the link SNR is highly stable, although the actual SNR values varies quickly. This is primarily attributed to the continuity of movement/rotation caused by objects or human that hold the VLC device. *Second*, VLC devices can readily support full-duplex, bi-directional transmission. To realize full-duplex, a conventional RF transceiver needs to cancel around 90dB of self-interference between its transmitting RF-chain to its own receiving RF-chain [4]. Yet for a VLC link, both the transmit (LED) and receive (PD) units are highly directional and can be easily isolated. Our experiments find negligible leakage interference between the LED and PD, even if they are co-located on the same device.

DLit leverages the above opportunities to realize a fine-grained in-frame rate adaptation scheme. It divides a packet into short-duration subframes such that channel varies negligibly therein. It *calibrates* the uplink and downlink, so that a transmitter can estimate its outgoing link quality directly through incoming signals' SNR levels. Instead of using packet-level SNR statistics [5], it *estimates* SNR at subframe level, without using extra preambles. Instead of using an averaging of historical SNR values, it establishes a Kalman filter that *predicts* the SNR of next subframe based on first- and second-order statistics. Finally, it *adapts* the modulation/coding level of next subframe on-the-fly.

Testbed implementation and experiments. To validate the principles behind DLit, we build a VLC software-radio testbed by extending the WARP platform [6] with third-party ADC/DAC and customized LED/PD drive circuits. We further prototype the PHY-layer modulation/coding in 802.15.7, and then implement DLit's adaptation protocol on top. Under realistic usage scenarios, our experiments demonstrate that DLit can accurately track the link dynamics and perform rate adaptation in a way that approximates an oracle. In contrast, conventional wireless rate adaptation protocols, either based on loss-statistics or packet-level SNR model, result in multi-folds lower throughput. To our knowledge, DLit represents the first system that explores the unique challenges/opportunities in VLC rate adaptation through testbed experiments.

The rest of the paper is structured as follows. Sec. II presents a background on VLC. Sec. III motivates DLit through a measurement study. We detail DLit's design components in Sec. IV and evaluate its performance in Sec. VI. Sec. VII discusses related work and finally, Sec. VIII concludes the paper.

II. Background

The IEEE 802.15.7 [7] has standardized a VLC system with three PHY layer modes, offering bit-rate range of 11.67 Kbps to 266.6 Kbps, 1.25 Mbps to 96 Mbps, and 12 Mbps to 96 Mbps, respectively. The PHY modes differ mainly in their sampling bandwidth, which can vary between 200 kHz and 120 MHz. Each mode operates either on-off keying (OOK) or variable pulse-position modulation (VPPM). A discrete set of bit-rate options can be created by controlling the redundancy level of an outer-layer error correction code, including Reed-Solomon and Convolutional Code. The optimal choice should balance the redundancy and error-protection capability, in order to maximize the link throughput. Such choice depends on channel quality and is the primary focus on our work.

802.15.7 assumes both sides of a link adopt LED/PD transceivers. This can be applicable between a ceiling-mounted light fixture and VLC-equipped autonomous devices, *e.g.*, robotic cleaner and medical devices [3]. For handheld devices, an LED transmitter may be power hungry and cause discomfort to human eyes. Therefore, alternative invisible medium such as WiFi and infrared has been proposed [1]. Our DLit scheme relies on a bi-directional link, where the downlink (from ceiling-mounted AP to mobile user device) is always using an LED transmitter, whereas for the uplink either LED or infrared can be used.

VLC strongly relies on a Line-of-Sight (LOS) communication path, although a NLOS path can be converted to LOS using steerable mirrors [8]. As we will show through experiments, VLC exhibits strong link dynamics even in an LOS channel.

III. Understanding VLC Link Dynamics

In this section, we conduct testbed experiments to investigate the spatial and temporal dynamics of a VLC link, which are critical factors for triggering the link-rate adaptation. We have built a software-radio platform (detailed in Sec. V) that supports 802.15.7-like VLC between an LED and PD, and between Infrared LED (IrLED) and PD (IrPD). Here we mainly examine how the VLC link quality varies in practical environment. Fig. 1 illustrates the key parameters that characterize our setup. We place a transmitter at a fixed location, facing a receiver that resides on a plane 1.5m away. This setup aims to emulate a real VLC usage case, where a ceiling-mounted AP serves a client (*e.g.*, mobile device held by a user) that can be arbitrarily located and moving within its coverage.

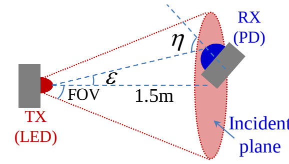

Fig. 1. Key parameter configurations of the VLC link in our testbed setup. ϵ : emission angle of TX; η: incidental angle *w.r.t.* RX; FOV: TX's field of view.

To estimate link quality, the transmitter sends a known OOK-modulated PN-sequence with length 64. After decoding the symbols, the receiver computes the post-processing SNR using a mature data-aided estimation algorithm (Sec. IV-C).

A. SNR Variation Over Space

We first examine the link quality when the PD varies its location within the LED's coverage area. Specifically, we measure the received SNR levels within a 1m$\times$1m region on the incident plane, sampled at 25 cm granularity. Fig. 2(a) plots the SNR distribution across the sampling points. We see that the SNR exhibits a smooth but steep degradation from the peak. At the center, the PD is straightly facing the LED with shortest link distance, resulting a high SNR of around 60 dB. Near the edge of the coverage area, the SNR plummets to near -5 dB, where the transmitted symbols are barely decodable. Interestingly, along each axis on the incident plane, the SNR curve roughly follows a one-cycle $\cos$ function, which is consistent with theoretical prediction of VLC channels in LOS [9].

Fig. 2(b) further illustrates a 2-D contour plot of the SNR distribution, where we represent the SNR of each 25cm$\times$25cm square by averaging the nearby 4 sampling points. We observe that the SNR levels again degrade sharply from the center, and is almost symmetry along all directions.

One may wonder if the invisible infrared link assumes different characteristics. To provide an empirical answer, we repeat the above experiments with a pair of IrLED and IrPD link. The results (Fig. 3) show that the infrared link has consistent spatial SNR distribution with the visible light, although the absolute SNR values differ slightly due to different hardware front-end. This provides us a first hint that the same link adaptation algorithm can apply for both visible light and infrared links.

B. Link Dynamics Over Relative Angle

In practical use of VLC, a handheld receiver may not always face straight towards the LED fixture. Human posture/gait change, hand shaking while walking, *etc.*, can all vary the incoming light signals' incident angle ϵ (defined in Fig. 1). To profile such effects, we place the RX directly under the TX

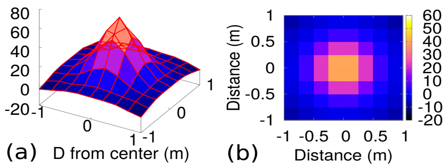

Fig. 2. VLC SNR distribution over position.

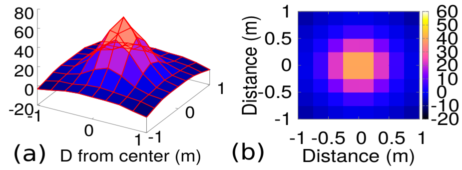

Fig. 3. Infrared SNR distribution over position.

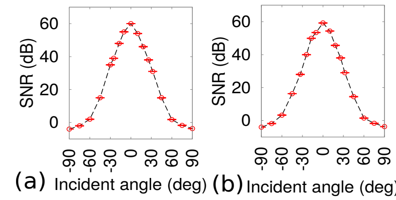

Fig. 4. SNR distribution over angle:(a) VLC (b) IR.

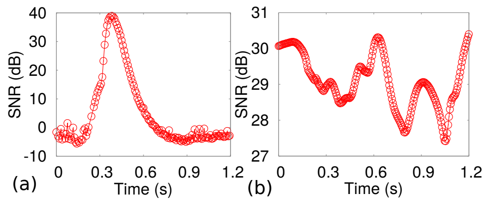

Fig. 5. SNR variation in walking in-and-out scenario:(a) VLC (b) WiFi

and then vary ϵ. Fig. 4(a) plots the resulting mean SNR and *std.* over 1 minute of continuous frame transmissions. We make two observations. *First*, the SNR again manifests as a smooth curve roughly following a $\cos$ curve over ϵ. It transits sharply from around 60 dB to -5dB as ϵ changes from $0°$ (PD directly facing LED) to $90°$ or $-90°$ (PD facing perpendicularly to LED). This implies that the VLC channel is highly directional and sensitive to angle variation between the TX and RX. *Second*, for each angle, the SNR varies negligibly over time. The *std.* is virtually 0 across all the experiments. Therefore, the VLC channel is highly deterministic for a given node location/orientation.

Fig. 4(b) shows the same experiments for the infrared link, which again exhibits similar characteristics as the visible light.

C. Link Dynamics in Mobile Environment

Based on the above observations, we expect that RX's mobility can create enormous link dynamics since the movement alters both its position and incident angle relative to the TX. We verify this hypothesis by imitating a practical scenario, where a user holds the RX and walks passing by the TX's coverage area. We intentionally face the RX up towards the ceiling and avoid shaking, such that the link dynamics are only induced by position change over time. As a contrast, we also repeat the experiment using a pair of WARP transmitter and receiver, running the 802.11g protocol at 2.412 GHz RF channel, and placed at the same locations as the TX and RX.

Fig. 5 plots the SNR variation sampled at a granularity of 20 ms. The VLC link experiences a drastic SNR change during the walk in-and-out: within 400 ms, the SNR raises from the noise floor (around -5 dB) to the peak level of 40 dB, and then back to noise floor. At the steepest region, it takes only 10 ms for the SNR to change by 10 dB. Considering the low rate of 802.15.7 especially in PHY mode I [7], the typical packet duration can be 10 ms for a 1.5KB packet. This means *the VLC channel already suffers from 10 dB of SNR change within one packet duration*. Apparently, the conventional per-packet rate adaptation mechanism no longer works here.

In contrast, the WiFi link exhibits a variation of less than 4 dB across the entire walk-in-and-out session. This is primarily because it is omni-directional and hence unaffected by the relative angle between transmitter and receiver. Although the TX-RX distances can vary, multipath reflections tend to "fill" the weak-signal regions, resulting in relatively consistent SNR.

An additional observation is that the *VLC link exhibits a predictable trend.*. In particular, as the SNR rises from the lowest to the highest level (or vice versa), the first-order statistics are roughly consistent. In contrast, although the long-term SNR variation of WiFi is small, the short-term variation is highly unpredictable due to sophisticated multipath reflections.

IV. DLit: Predictive In-frame Rate Adaptation

DLit is a fine-grained rate adaptation scheme that leverages the full-duplex and predictability features of VLC links to tackle link-quality variations within a frame. Fig. 6 illustrates the system architecture of a DLit node. Given a MAC-layer packet (frame), DLit divides it into subframes whose length falls well within the channel coherence time. DLit transmitter estimates its forward-link quality by inspecting the concurrent reverse-link, and then predicts the optimal bit-rate on a per-subframe basis. More specifically, its adaptation framework comprises the following modules:

(i) Full-duplex calibration, a one-time execution that ensures channel reciprocity between the bi-directional links, allowing DLit to obtain the forward-link SNR directly from the reverse-link, instead of per-subframe feedback which incurs significant overhead and latency.

(ii) Blind in-frame SNR estimation, which directly uses the order-statistics of received symbols to estimate the reverse-link SNR, instead of using preambles which wastes channel time.

(iii) Subframe SNR prediction that leverages the stability of the VLC channel's first and second-order statistics (Sec. III), and uses a Kalman filter to predict the SNR level of next subframe. The predicted SNR is then mapped to an optimal bit-rate, which guides the modulator to encode the next subframe and send it through the LED.

One may wonder if the full-duplex concurrent transmission requirement limits DLit's applicability. We argue that without full-duplex in-frame SNR estimation, a VLC link needs to either remain at the most conservative MCS, or rely on per-packet feedback which can no longer handle link dynamics (Sec. VI). Thus, it is worthy for the reverse link to transmit concurrently. If no packets are queued up, it can transmit known data to support the forward-link rate adaptation.

In what follows, we detail each of DLit's components and corroborate them with micro-benchmark experiments.

A. Subframe Structure

1) Frame Format: Fig. 7 illustrates DLit's frame structure. Taking the downlink transmission (AP→client) as an example. The AP prepends each data frame with an 802.15.7-compatible preamble comprised of a sequence of 16 OOK-modulated 1's, used for frame-level synchronization between the client and the

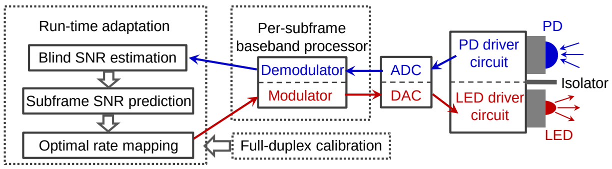

Fig. 6. Architecture of a VLC node running DLit.

AP. Following the preamble is a sequence of data subframes with *adaptively installed MCS headers*.

The initial subframes always start with the lowest modulation and coding (MCS) level. Once the AP obtains data from the reverse uplink, it predicts the SNR for the next downlink subframe period and chooses the throughput-optimal MCS. If the MCS differs from the current one, it prepends the next subframe with an *MCS header* that indicates the new MCS.

As shown in Fig. 7, the MCS header comprises 24 bits modulated with the lowest bit-rate. The first 20 bits are filled with 1's to distinguish it from normal data, followed by a 4-bit index of the new MCS level. The client needs to first decode the leading 24 bits of each subframe with the lowest MCS. Upon detecting an MCS header, it proceeds to decode the data part with the new MCS. Otherwise, it rolls back to the beginning of the subframe and decodes it with the previous MCS.

2) Subframe Size and Pre-modulation: DLit uses a fixed subframe size in terms of the number of bits instead of duration. The rationale behind this design choice is to simplify the baseband modulation. Given a fixed number of bits per subframe, the modulator can *pre-modulate* the packet data using all candidate MCS levels before transmission starts. The premodulated raw signals can be cached in a buffer. During frame transmission, the signals of each subframe can be directly fetched from the buffer, instead of remodulated at runtime which may incur nontrivial latency and disrupt the transmission.

To facilitate pre-modulation, the duration of each subframe should be shorter than the channel coherence time. DLit satisfies this by considering the worst case: Suppose the modulator uses the lowest MCS with bit-rate R, and the shortest time when channel remains stable is t_0. Then the subframe size is fixed to $t_0 R$ bits. We empirically set t_0 to 2 ms which, according to our measurement, is the shortest time within which the channel can change by up to 2 dB in indoor walking scenarios. In our DLit implementation, the lowest MCS level is 500 Kbps, and thus the subframe size should be below 130 bytes. We fix it to a conservative value of 100 B.

B. Full-duplex VLC Link and Reciprocity Calibration

In this section, we first verify the feasibility and practical performance of a full-duplex VLC node, and then describe the reciprocity calibration mechanism in DLit that leverages a semi-symmetric property of bi-directional VLC links.

1) Self Interference in Full-duplex VLC: Both the LED and PD have close-to-zero response at $90°$ emission/incident angle. Thus, when the LED is placed on the same plane and pointing to the same direction as the PD, the self-interference effect should be minimal.

To verify this intuition, we measure the SNR degradation when a PD (IrPD) receiving module is interfered by a LED

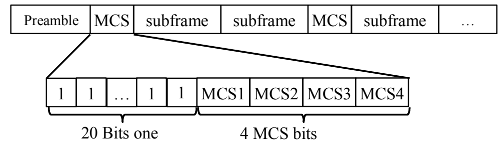

Fig. 7. Frame structure of DLit.

(IrLED) transmission module co-located on the same node. An opaque plastic paper, with the same height as the LED, is placed in between them to serve as an isolator (Fig. 6). This setup emulates a real VLC mobile device with the LED and PD embedded on the front-panel (similar to the infrared emitter and sensor in modern smartphones).

Fig. 8 shows the impact of self-interference. For a bi-directional VLC link, the SNR loss is always below 0.3 dB even if the LED and PD are placed 2.5cm apart. Thus, the self-interference is negligible. Moreover, the VLC TX (LED) imposes almost zero interference to the infrared RX (IrPD) since the IrPD has a black optical lens that filters out the visible light. But the VLC PD has no filter lens, and hence the infrared TX (IrLED) still generates interferes with the VLC RX, albeit at a negligible level of below 0.23 dB.

To summarize, *in a VLC network with a visible-light link, no matter if the other link is infrared or visible-light, the mutual-interference between the bi-directional links is negligible, and full-duplex can be supported even without sophisticated self-interference cancellation as required by RF networks [4].*

2) Semi-symmetric Channel Response in Full-duplex VLC: The VLC channel gain comes from two factors: *(i)* the optical pathloss that only depends on the distance/angle between the LED and PD. The pathloss factor may vary but is symmetric especially for the directional LOS channel [9]. *(ii)* the hardware-induced gain that depends on FOV, transmit power, RX amplifier gain, *etc.*, which is typically asymmetric but stable over time. To examine such factors, we measure the SNR of bi-directional links between a static AP and mobile client. Fig. 9(a) snapshots the SNR variation over time, when both links use the same VLC hardware. We see that the SNR curves are fully overlapping, indicating symmetric full-duplex channels.

When the reverse link is replaced with infrared with the same FOV (Fig. 9(b)), the two links' SNR differs by only a constant. This is mainly because of a constant hardware gain difference in the transceiver circuits, which can be easily calibrated.

We then use two PDs with different FOVs ($20°$ vs. $45°$) for the bi-directional links. The resulting channel gain difference is no longer a constant due to the different angular responses, as shown in the SNR variation (Fig. 9(c)).

Fortunately, we observe that the typical angular response of PD or LED is almost stable within the range of FOV, and falls sharply to 0 near the boundary [10]. This is also manifested in Fig. 9(c) — the gain difference between the two links becomes abnormal only in those regions where SNR is extremely low (due to angle misalignment), where connectivity is barely guaranteed. Therefore, in the common cases, we can safely approximate the gain difference and compensate for it using a constant value representing the high-SNR case.

3) Calibration of the Full-duplex SNR Sensing: Since the SNR difference between forward and reverse link is only due

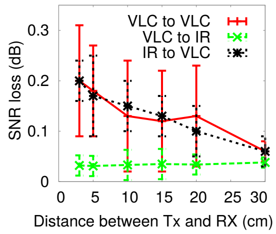

Fig. 8. SNR loss by self-interference

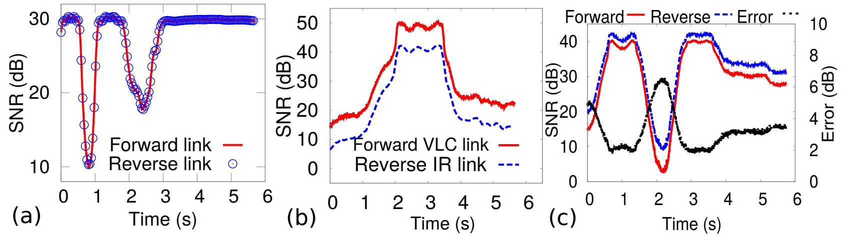

Fig. 9. SNR traces in different hardware settings : (a) same hardware; (b) VLC and infrared; (c) different FOVs.

to time-invariant hardware properties and can be approximated as a constant, the AP and client can conduct a one-time-calibration to compensate for the gain difference. During the association procedure, the client should be placed within the AP's FOV, exchange short subframes and the corresponding SNR values with it, and keep using the SNR difference for gain compensation afterwards. Notably, the client's gain difference can be shared among all APs to save the calibration overhead.

In the actual implementation of DLit, we choose to compensate the gain difference using a slightly small value than the measured common case. This may cause underselection of MCS, but is less harmful to system performance than overselecting that leads to packet losses (Sec. VI).

To verify the accuracy of the full-duplex calibration, we measure the calibration error, *i.e.*, difference between the SNR of a VLC forward-link and the SNR estimated from a calibrated infrared reverse-link with FOV. Fig. 10 plots the error CDF across 10^4 frames, where we intentionally move and shake the client to create a variety of SNR scenarios. We can see even with an unusually large FOV difference of $45°$, we can keep the 90-percentile calibration error below 1.5 dB. With FOV difference below $20°$, the median error is below 0.5 dB.

C. In-Frame Blind SNR Estimation

DLit employs a blind stochastic estimator to obtain the reverse-link SNR without adding extra preamble to every subframe. Unlike conventional preamble-based SNR estimator [11], the blind estimator works even with a few unknown symbols inside each short subframe.

More specifically, observing the simplicity of the OOK modulation in VLC, we expect the signal strength statistics of the raw samples are sufficient for SNR estimation. It is known that in an AWGN channel, for modulation schemes with a discrete set of signal magnitudes (*e.g.*, as in OOK), the second and fourth order moments of signal strength can approximate the SNR as:

$$\mathrm{SNR}_{\mathrm{M2M4}} = \frac{\frac{1}{2}\sqrt{6M_2^2 - 2M_4}}{M_2 - \frac{1}{2}\sqrt{6M_2^2 - 2M_4}}, \tag{1}$$

In DLit, a node uses such an *M2M4 estimator* to directly obtain SNR estimation from raw samples without knowing the corresponding digital symbols. Denote y_n as a sequence of OOK samples inside one subframe, and N_{sym} is the size of SNR observation window which is set to half of the subframe size by default. Then the moments can be represented as:

$$M_2 \approx \frac{1}{N_{sym}} \sum_{n=0}^{N_{sym}-1} |y_n|^2, \text{ and } M_4 \approx \frac{1}{N_{sym}} \sum_{n=0}^{N_{sym}-1} |y_n|^4$$

To verify the accuracy of the estimator, we compare it with an oracle data-aided estimator, which knows all the OOK symbols (as if they are all known preambles) in a subframe and can directly compute SNR from the mean signal power and noise variance. Fig.11(a) plots the estimation error under a variety of SNR levels. Error bars represent the *std.* over 10^4 frames. We can see the maximum estimation error is only around 0.4dB with small variations.

Fig. 11(b) further shows the CDF of estimation error across 3 experiments where the client walks by/below the AP. Each experiment takes around 5 seconds. We again see a small median error below 0.1 dB. From both sets of experiments, we can conclude that the blind M2M4 estimator has a small error that falls well below the SNR gap between two neighboring MCS levels (typically 2 to 5 dB, as will be shown later), and therefore it is well suited for in-frame SNR estimation.

D. Subframe SNR Prediction

From the experiments in Sec. III, we saw that the VLC link SNR is highly predictable, because it is dominated by a directional LOS path. The main factor that varies the SNR is device movement or rotation. As such dynamics are smooth and roughly linear in a short-term, we can use a discrete-time linear dynamic system to approximate the SNR evolution over time. Specific to DLit, we use a second-order Kalman filter to predict the SNR for each next subframe, as described below.

1) Dynamic Model of SNR Variation Process: We first establish a dynamic system model for the SNR variation process. We discretize the time using the index of subframes, denoted as t. Let x_t be the state space at time (subframe) t, then the corresponding dynamic model can be represented by:

$$x_{t+1} = Fx_t + w_t, \tag{2}$$

which is referred to as the *state equation* of the Kalman filter model. Here F is the state transition model applied to the discrete process x_t. w_t is the process noise that follows the Gaussian distribution with covariance Q, *i.e.*, $w_t \sim \mathcal{N}(0, Q)$.

In our design, we assume the acceleration (second order derivative) of the SNR variation process, denoted as a_t, is stable between time slot t and $t+1$. We further assume a_t is normally distributed with mean 0 and standard deviation σ_a. The state space of the process x_t is described as $x_t = (\begin{array}{cc} \rho_t & \rho'_t \end{array})^{\mathrm{T}}$, where x_t is the SNR and x'_t the velocity of SNR variation. Based on the second law of motion, we can describe x_t as: $x_{t+1} = Fx_t + Ga_t$.

$$F = \begin{pmatrix} 1 & \Delta t \\ 0 & 1 \end{pmatrix} = \begin{pmatrix} 1 & 1 \\ 0 & 1 \end{pmatrix}, G = \begin{pmatrix} \frac{\Delta t^2}{2} \\ \Delta t \end{pmatrix} = \begin{pmatrix} 0.5 \\ 1 \end{pmatrix}$$

Recall the system is discrete and thus $\Delta t = 1$. Back to the state equation (2), we have $w_t = Ga_t$ and its covariance

$$Q = GG^T\sigma_a^2 = \begin{pmatrix} 0.25 & 0.5 \\ 0.5 & 1 \end{pmatrix}\sigma_a^2 \tag{3}$$

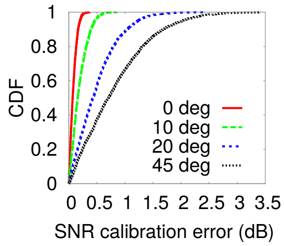

Fig. 10. SNR calibration error.

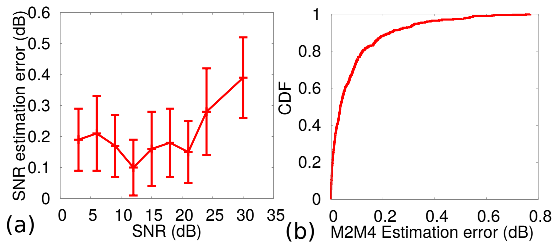

Fig. 11. (a) M2M4 estimation error vs SNR; (b) CDF of estimation error.

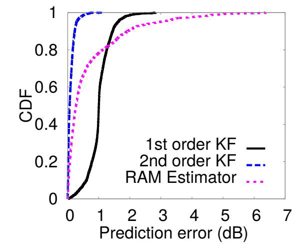

Fig. 12. Kalman prediction error.

It is hard to measure the acceleration of the SNR variation process since the movement pattern is not stable. In DLit, the variance of acceleration is empirically set to a relatively large value $\sigma_a^2 = 2$ to feed more uncertainty into the Kalman filter to enable fast adaptation when link SNR changes suddenly.

To predict the state of next time slot, Kalman filter uses a run-time measurement-driven self-calibration algorithm. Specific to DLit, we define a *measurement equation* as:

$$z_t = Hx_t + v_t \tag{4}$$

where z_t denotes the observation variable, *i.e.*, the calibrated SNR of the forward-link in the current subframe, measured via the backward-link (Sec. IV-B). Denote R as the covariance of z_t. H represents the observation model which maps the true SNR state space to measured state space and v_t is the *measurement noise* and $v_t \sim \mathcal{N}(0, \sigma_v^2)$.

Since only the SNR (and not its velocity) can be directly measured, we have $H = \begin{pmatrix} 1 & 0 \end{pmatrix}$ and $R = E\left[v_t v_t^{\mathbf{T}}\right] = \sigma_v^2$. We empirically set σ_v to 1 to model our SNR measurement and estimation error, which is small even in the worst case.

2) Kalman Filter Algorithm: Kalman filter runs two estimations in parallel: priori estimation and posteriori estimation. At time t, it performs priori estimation $x_{t+1|t}$ to predict the state of for next time slot $t+1$. It also performs posteriori estimation $x_{t|t}$ that combines both the current observation z_t and the priori estimation done at time $t-1$, *i.e.*, $x_{t|t-1}$. We use $P_{t|t}$ and $P_{t+1|t}$ to represent the covariance of $x_{t|t}$ and $x_{t+1|t}$.

In DLit, the Kalman filter has two phases at each time slot (subframe) t: *updating* and *predicting*. In the updating phase we calculate the posteriori estimation $x_{t|t}$ and covariance $P_{t|t}$, based on the SNR measurement z_t, as well as the $x_{t|t-1}$ and $P_{t|t-1}$, which are the priori estimation and covariance calculated at $t-1$. More specifically:

$$x_{t|t} = x_{t|t-1} + K_t\left(z_t - Hx_{t|t-1}\right) \tag{5}$$

$$P_{t|t} = P_{t|t-1} - P_{t|t-1}K_t H \tag{6}$$

where K_t is the *Kalman gain* matrix, defined as:

$$K_t = \frac{P_{t|t-1}H^{\mathrm{T}}}{HP_{t|t-1}H^{\mathrm{T}} + R} \tag{7}$$

Intuitively, K_t depends on the relative magnitudes of matrix R and $P_{t|t-1}$. The update equations (5) and (6) imply that, when the magnitude of R is small (*i.e.*, the measurements are accurate), the state prediction depends mostly on the measurements z_t. When the state was predicted accurately, then $HP_{t|t-1}H^{\mathrm{T}}$ is small compared to R, and the filter mostly ignores the measurements z_t, relying instead on the prediction derived from the priori state $x_{t|t-1}$. It is known that the form of (7) makes this Kalman filter an MMSE (minimum mean square error) predictor [12].

In the *predicting* phase, we calculate the priori estimation $x_{t+1|t}$ and $P_{t+1|t}$ based on the posteriori estimation $x_{t|t}$ and $P_{t|t}$ and the dynamic system model in (2), as follows:

$$x_{t+1|t} = Fx_{t|t} \tag{8}$$

$$P_{t+1|t} = FP_{t|t}F^{\mathrm{T}} + Q \tag{9}$$

We then use $\rho_{t+1|t}$, the first row of $x_{t+1|t}$, as a prediction of the SNR of the next subframe $t+1$, and subsequently select the optimal MCS (Sec. IV-E).

3) SNR prediction error: To verify the Kalman filter design, we implement the prediction algorithm and run it over a real SNR trace created by moving and rotating a mobile VLC client under the FOV of an AP fixture. For benchmark comparison, we also implement a first-order Kalman filter that assume the first-order state transition (velocity of SNR variation) is stable between subframes, and also a typical moving-average predictor called RAM [5] which is designed for 802.11 packet-level rate adaptation. The trace is collected using 800us frame with a 3.5ms frame interval using our testbed (Sec. V).

Fig. 13 plots the real SNR and one step prediction along with the evolution of prediction error over time. Across the entire process, the Kalman model's prediction error is below 0.6 dB, which falls well below the SNR gap between two MCS levels and is unlikely to affect the rate selection significantly. The slight prediction error is mainly because the Kalman filter uses a linear model, but the distribution over space or incidence angle is nonlinear (Sec. III). The first-order Kalman filter exhibits errors mostly below 1.6 dB. The error concentrates on the transition points due to sudden device movement or rotation, where the first-order predictor has a lagging effect. The RAM predictor causes even larger errors at such points, mainly because it assumes the SNR is stable across frames, and thus the next frame's SNR can be predicted by averaging historical SNR, which no longer holds for VLC links.

E. Mapping from SNR to MCS

To select the optimal MCS for a given SNR prediction, we consider both the MCS bit-rate and its subframe delivery ratio (FDR). Recall DLit's subframe size is fixed to 100 bytes. A fixed frame size allows us to use a fixed table to map SNR to FDR. To generate the mapping table, we again use our VLC testbed to generate different link SNR levels, and measure the resulting average FDR across 10^4 subframes. Fig.14 shows the resulting SNR-to-FDR mapping for each MCS level, which provides a mapping function $\mathrm{FDR}(\mathrm{MCS}, \mathrm{SNR})$.

Let $R(\mathrm{MCS})$ be the bit-rate corresponding to one MCS level, then the throughput-optimal MCS selection MCS^* is given by

$$\mathrm{MCS}^* = \underset{\mathrm{MCS}}{\mathrm{argmax}} \quad \mathrm{FDR}(\mathrm{MCS}, \mathrm{SNR}) \times R(\mathrm{MCS}) \tag{10}$$

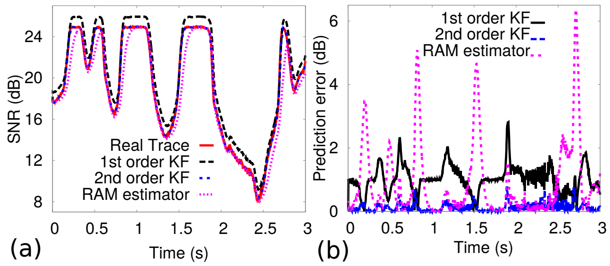

Fig. 13. (a) Trace of real SNR and one-step lookahead prediction result; (b) The prediction error.

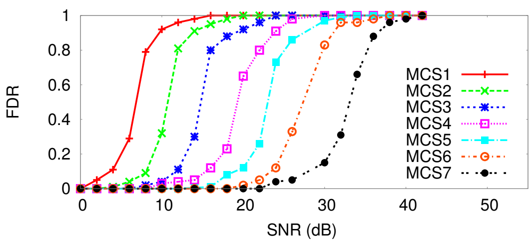

Fig. 14. SNR to FDR mapping

V. Implementation

To our knowledge, there is no off-the-shelf platform that supports software-defined implementation of rate adaptation algorithms for 802.15.7-like VLC systems. So we build a VLC software-radio by extending the WARP [6], a reconfigurable baseband FPGA processor originally designed for RF communications. Fig. 15 illustrates our software-radio platform with system diagram shown in Fig. 16.

At the transmitter front-end, we use 3W plug-and-play LED with a variety of FOVs. We built an AC power-amplifier circuit from scratch to drive the LED. A Bias-Tee module is used to combine the modulated AC waveform and the DC power that lights up the LED to meet its current/voltage requirements. The AC waveform is generated by modulating digital bits through a software-defined modulation module on WARP, passing a baseband filter, and then converted to analog signals by a third-party ADC board (FMC150) that we added on WARP.

At the receiver side, the PD's received signals are amplified, injected to the FMC150's ADC module, and consequently converted into real-valued digital samples that can be processed by the software-defined demodulator. Both the PD and LED can be replaced by IrPD and IrLED with similar power requirement.

We developed an FPGA driver program that allows WARP to interface with the FMC150 and feed digital samples to or receiver from it. The WARP board is connected to a host PC that runs the modulation/demodulation modules that we developed on top of the WARPLab driver library.

We implement the OOK modulation/demodulation schemes along with DLit's subframe construction mechanism (Sec. IV-A). Each sequence of subframes follow a 802.15.7 preamble that acts as start-of-frame delimiter for TX-RX synchronization. Given a certain sampling bandwidth, 802.15.7's PHY mode only supports up to 3 modulation levels [7]. Our implementation fixes the sampling bandwidth to 8 MHz, but varies the outer-layer Reed-Solomn coding rate, to create up to 7 MCS levels, supporting bit rate from 500 Kbps to 4 Mbps.

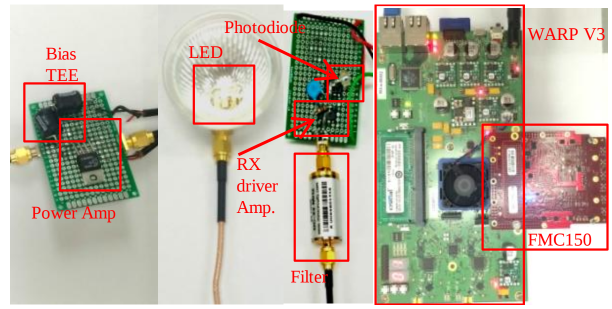

Fig. 15. VLC software-radio platform.

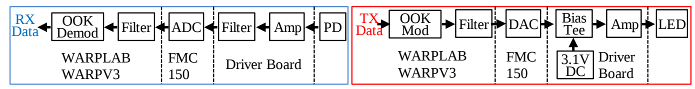

Fig. 16. Architecture of the VLC software-radio prototype

VI. Performance Evaluation

A. Experimental Setup

We evaluate the efficacy of DLit in a VLC network containing a ceiling-mounted AP and up to 3 mobile clients. Each AP-client pair is a full-duplex link with visible-light uplink and infrared downlink. Without loss of generality, we mainly focus on the rate adaptation performance at the client side that runs the aforementioned DLit implementation.

For comparison, we also port conventional 802.11 rate adaptation protocols which can be broadly classified into history-based and SNR-model-based approach [13]. The *former* uses historical packet loss statistics as an adaptation metric, and balances exploration and exploitation by occasionally upgrading the MCS level. We implement the RRAA protocol [13] for comparison, but remove the RTS/CTS related adaptation which is not supported in 802.15.7. The *latter* category uses per-packet SNR statistics to adapt the MCS level, wherein the SNR level is piggy-backed by the ACK from reverse link. We implement the SGRA [5] as a representative SNR-guided protocol. Besides, we develop an *oracle* rate adaptation algorithm that runs offline and directly computes the per-subframe optimal MCS based on the measured SNR values.

When running DLit on our testbed, we find the interface and signal demodulation latency of the WARP (plus PC host) can be several hundred milliseconds per-packet, which is unsuitable for fine-time adaptation. We circumvent this limitation by allowing both sides of the full-duplex link to continuously send a pre-built packet. The resulting inter-packet latency is reduced to 3–4 ms, with frame duration up to 800 μs. Each node first stores all the received *raw signal samples*, and then processes the samples, following the DLit design (Sec. IV) and benchmark protocol specifications. On top of the signal processing engine, we emulate a simple MAC with a virtual timer that simulates 802.15.7's random backoff, inter-frame time, and ACK.

B. Experimental Results

We evaluate the rate adaptation protocols on the mobile client, held by a user and subject to practical movement (walking) and shaking effects. We start with a default configuration of saturated UDP-like traffic with 1.5 KB packet size.

1) *Micro Benchmark:*

Responsiveness. Considering the highly dynamic VLC channel, we first examine how DLit responds to fast-changing link conditions. Fig. 17 plots a 3-second trace of SNR variation when the client is rotated randomly, and the corresponding MCS levels selected by different protocols over time. We observe that DLit follows the oracle protocol closely with only occasional lagging, owing to its fine-grained subframe-level rate adaptation. RRAA's exploitation strategy leads it to frequent underselection and overselection, mainly because the latency of exploitation results in outdated decisions. With the moving average prediction, SGRA can roughly follow the oracle's trend and shows small smaller variation, but the lagging effects become severe whenever the channel changes suddenly.

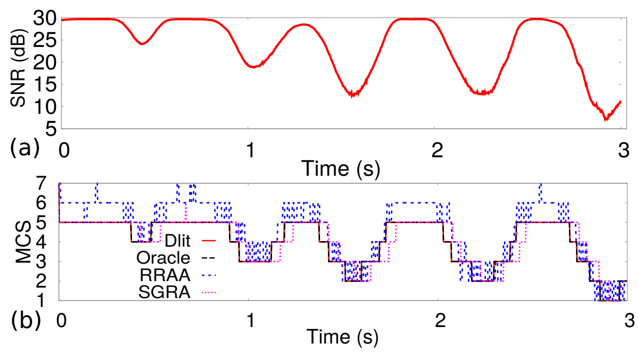

Fig. 17. SNR variation during RX shaking (rotation), and the corresponding response of different rate adaptation protocols.

Slow walking scenario. Fig. 18(a) plots the percentage of successful or failed transmission attempts across different MCS levels during a 5-minute slow-walking (0.5 m/s) scenario. Note that DLit and the oracle show similar distribution of transmission attempts, and around 10% of frame loss rate. The main reason behind frame loss is that the optimal MCS selection scheme aims for the highest throughput but not necessarily the highest frame delivery ratio (Eq. (10)). RRAA exhibits a very high frame loss rate at MCS 5 since it often switches to higher MCS to exploit the feasibility at the cost of transmission failure. Different from RRAA, the reason behind SGRA's transmission failure is the feedback latency which outdates the rate selection.

Fast walking scenario. When the client is walking fast at around 2 m/s, DLit introduces around 5% more frame losses than the oracle due to imperfect adaptation. RRAA and SGRA tend to have relatively more attempts on higher MCS levels (6 and 7) which leads to more transmission failures.

Fast walking and strong shaking scenario. We further emulate an extreme condition where the user walks fast while shaking the phone fiercely. From the results in Fig. 18(c), we see that DLit underselects the MCS in high-SNR region (high MCS indices). This is primarily because we use forward and reverse links with FOV difference of 25°, and the full-duplex calibration procedure tends to underestimate the forward link SNR because of the conservative full-duplex calibration (Sec. IV). However, the impact of underselection is much less compared with outdated decision making and overselection, which causes significant transmission failure as shown in the RRAA and SGRA cases.

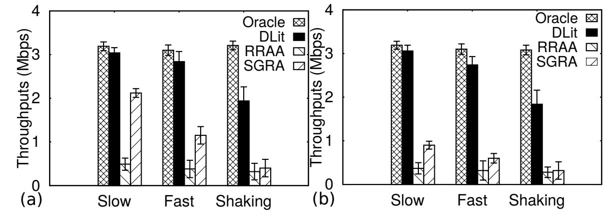

Fig. 20. Throughputs with different frame size: (a) 100B, (b) 1500B

2) *System Throughput Test:*

Throughput under different frame sizes. We evaluate the overall throughput performance under similar scenarios as above. Fig. 20 show the mean and *std.* of throughput across 10 repetitions for each scenario. We observe that frame size does not affect DLit's performance in a noticeable way, except in the shaking case, where the frequent use of MCS headers incurs more overhead and degrades the throughput by around 5%. In both the slow and fast walking scenarios, DLit's throughput is 90% to 96% close to the oracle. With 100B packets, DLit achieves around 4.3× over RRAA, and 2.2× over SGRA; with a larger packet size of 1.5KB, the gain escalates to around 5× and 3.4×, respectively. Longer frame tends to suffer more from the in-frame channel change, which cannot be handled by the per-packet feedback schemes. For the same reason, DLit outperforms both schemes by around 4× in the unusually fierce device shaking scenario, although its imperfectness compare with the oracle is also magnified.

Throughput with multiple clients. In this experiment, we repeat the fast-walking scenario, but vary the number of contending clients in the VLC network that runs the aforementioned 802.15.7 random access MAC. All clients use a packet size of 100 B. From the results (Fig. 19), we can see DLit achieves near-optimal throughput owing to its in-frame adaptation mechanism. SGRA shows reasonable performance in single-user case. But as the user population grows, the increasing inter-packet delay stales the SNR feedback, causing a super-linear drop of throughput. Similar consequence happens for RRAA. With only 4 users, DLit achieves 3.2× and 7.8× per-user throughput compared with these two schemes.

VII. Related Work

VLC: communications and application research. Nakagawa *et al.* pioneered the early research on VLC [14] with channel modeling and simulation validation. The majority of follow-on research focused on PHY-layer schemes to boost the VLC bit-rate. Besides the OOK/VPPM in 802.15.7, more efficient modulation schemes like OFDM have been proposed [15]. Combined with advances in photonic hardware (*e.g.*, laser LED/PD), multi-Gbps VLC links becomes feasible [16].

However, little work has discussed the adaptation between different rates for VLC. One recent system, VMRA [17], establishes parallel (MIMO) links between multiple LEDs and a multi-pixel camera image sensor. It adaptively triggers a subset of links according to historical packet BER feedback. In DLit, we have observed that packet-level adaptation cannot handle device mobility. Thus, DLit employs a predictive rate adaptation mechanism, facilitated by calibrated full-duplex SNR estimation and in-frame rate adaptation. We believe these mechanisms

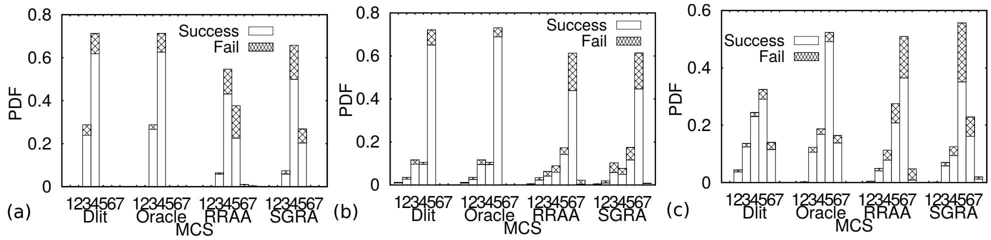

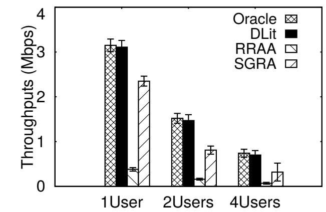

Fig. 18. PDF of MCS in different scenario (a) Slow walking (b)Fast walking (c) fast walking with strong shaking. Fig. 19. Throughput with multiple clients.

can be equally applied to boost mobile visual-MIMO links.

Besides dedicated hardware, VLC has shown to be feasible for legacy mobile devices. For example, COBRA [18] and LightSync [19] prototype short-range point-to-point communication links through screen-camera pairs. Achievable rate of such systems largely depends on alignment (image distortion) between the transceivers, and thus the rate adaptation may require a more complicated model that extends DLit.

Rate adaptation in wireless networks. Wireless rate adaptation protocols have been extensively studied especially for 802.11 networks. Based on the channel-quality metric, they can be classified into frame-loss based and SNR based approaches (see, *e.g.*, [5], [13], [20]). Both need at least one historical packet transmission to access the achievable rate for the next packet. Recent cross-layer designs have pushed the limit of wireless rate adaptation through new SNR metrics [21] specifically applied to 802.11 OFDM modulation.

For a VLC link, our measurement has shown that SNR changes significantly even within one frame (Sec. III), which invalidates the assumption behind wireless rate adaptation protocols. DLit solves this problem through two techniques that are uniquely applicable to VLC: predictive preamble-free SNR estimation and subframe-level rate adaptation. Fine-grained subframe processing is recently leveraged in Micro-ACK [22] to realize in-frame retransmission. Yet to our knowledge, DLit is the first work that achieves subframe-level rate adaptation leveraging the predictable and full-duplex VLC links.

We note that the first version of IEEE 802.11 incorporated infrared communications which was later replaced by RF technologies. Early work on infrared discussed link rate adaptation, but only based on theoretical channel models [23], [24]. As we have shown experimentally, infrared and visible-light links show similar response to network dynamics, and thus DLit's practical solution is applicable to both.

VIII. Conclusion

We have presented DLit, a fine-grained in-frame rate adaptation scheme uniquely designed for VLC networks and verified through testbed experiments. The most valuable lesson we learned from DLit is the fundamentally new challenge (highly dynamic SNR variation within one packet), as well as new opportunities (channel predictability and full-duplex) for a rate-adaptive VLC link. The insights from DLit can foster new network protocol designs that push the vision of ubiquitous indoor VLC networking. We remark that DLit only deals with the rate adaptation problem for a connected VLC link. Maintaining connectivity under blockage, shadowing, handoff, *etc.*, is a matter of our future work.

Acknowledgement

The work reported in this paper was supported in part by the US NSF under Grant CNS-1318292, CNS-1343363, CNS-1350039, CNS-1404613 and National High Tech. RD Program under Grant 2014AA01A706 in China.

References

[1] H. Elgala, R. Mesleh, and H. Haas, "Indoor Optical Wireless Communication: Potential and State-of-the-Art," *IEEE Communications Magazine*, vol. 49, no. 9, 2011.

[2] L. Grobe, A. Paraskevopoulos, J. Hilt, D. Schulz, F. Lassak, F. Hartlieb, C. Kottke, V. Jungnickel, and K.-D. Langer, "High-speed Visible Light Communication Systems," *IEEE Comm. Magazine*, 2013.

[3] R. Murai, T. Sakai, H. Kawano, Y. Matsukawa, Y. Kitano, Y. Honda, and K. Campbell, "A Novel Visible Light Communication System for Enhanced Control of Autonomous Delivery Robots in a Hospital," in *IEEE International Symposium on System Integration*, 2012.

[4] D. Bharadia, E. McMilin, and S. Katti, "Full Duplex Radios," in *ACM SIGCOMM*, 2013.

[5] J. Zhang, K. Tan, J. Zhao, H. Wu, and Y. Zhang, "A Practical SNR-Guided Rate Adaptation," in *Proc. of IEEE INFOCOM*, 2008.

[6] "WARP Project," 2014. [Online]. Available: http://warpproject.org

[7] S. Rajagopal, R. Roberts, and S.-K. Lim, "IEEE 802.15.7 Visible Light Communication: Modulation Schemes and Dimming Support," *IEEE Communications Magazine*, vol. 50, no. 3, 2012.

[8] N. Hamedazimi, H. Gupta, V. Sekar, and S. R. Das, "Patch Panels in the Sky: A Case for Free-space Optics in Data Centers," in *HotNets*, 2013.

[9] Z. Ghassemlooy, W. Popoola, and S. Rajbhandari, *Optical Wireless Communications: System and Channel Modelling with MATLAB*. CRC Press, 2012.

[10] K. Cui, "Physical Layer Characteristics and Techniques for Visible Light Communications," Ph.D. dissertation, UC-Riverside, 2012.

[11] D. Pauluzzi and N. Beaulieu, "A comparison of snr estimation techniques for the awgn channel," *Communications, IEEE Transactions on*, vol. 48, pp. 1681–1691, 2000.

[12] G. Welch and G. Bishop, "An introduction to the kalman filter," 1995.

[13] S. H. Y. Wong, H. Yang, S. Lu, and V. Bharghavan, "Robust Rate Adaptation for 802.11 Wireless Networks," in *ACM MobiCom*, 2006.

[14] T. Komine and M. Nakagawa, "Fundamental Analysis for Visible-Light Communication System Using LED Lights," *IEEE Transactions on Consumer Electronics*, vol. 50, no. 1, 2004.

[15] H. Elgala, R. Mesleh, and H. Haas, "Indoor Broadcasting Via White LEDs and OFDM," *IEEE Trans. on Consumer Electronics*, vol. 55, no. 3, 2009.

[16] M. Noshad and M. Brandt-Pearce, "Can Visible Light Communications Provide Gb/s Service?" *CoRR*, vol. abs/1308.3217, 2013.

[17] A. Ashok, M. Gruteser, N. Mandayam, T. Kwon, W. Yuan, M. Varga, and K. Dana, "Rate Adaptation in Visual MIMO," in *IEEE SECON*, 2011.

[18] T. Hao, R. Zhou, and G. Xing, "COBRA: Color Barcode Streaming for Smartphone Systems," in *Proc. of ACM MobiSys*, 2012.

[19] W. Hu, H. Gu, and Q. Pu, "LightSync: Unsynchronized Visual Communication over Screen-camera Links," in *Proc. of ACM MobiCom*, 2013.

[20] X. Chen, P. Gangwal, and D. Qiao, "Practical Rate Adaptation in Mobile Environments," in *Proc. of IEEE PerCom*, 2009.

[21] D. Halperin, W. Hu, A. Sheth, and D. Wetherall, "Predictable 802.11 Packet Delivery from Wireless Channel Measurements," in *Proc. of ACM SIGCOMM*, 2010.

[22] J. Zhang, H. Shen, K. Tan, R. Chandra, Y. Zhang, and Q. Zhang, "Frame Retransmissions Considered Harmful: Improving Spectrum Efficiency Using Micro-ACKs," in *Proc. of ACM MobiCom*, 2012.

[23] L. Diana and J. Kahn, "Rate-Adaptive Modulation Techniques for Infrared Wireless Communications," in *Proc. of IEEE ICC*, 1999.

[24] J. Grubor, V. Jungnickel, and K.-D. Langer, "Adaptive-Modulation Technique in Wireless Infrared Indoor Communications," in *ITG Symposium on Photonic Networks*, 2006.