LiRa: a WLAN architecture for Visible Light Communication with a Wi-Fi uplink

Sharan Naribole, Shuqing Chen, Ethan Heng, and Edward Knightly
ECE Department, Rice University

***Abstract*—Visible Light Communication (VLC) can dual purpose energy efficient LED-based lighting infrastructure for both illumination and data communication. Unfortunately, this dual-purposing is only inherent in the downlink direction, from infrastructure illumination sources to mobile devices. In this paper, we design, analyze, and implement LiRa, a Light-Radio WLAN that fuses light and radio capabilities in an integrated system design without requiring mobile devices to emit light or infrared. We design an uplink control channel for LiRa that is Wi-Fi compliant, has a controllable impact on airtime taken from legacy Wi-Fi clients, and efficiently scales with increasing VLC user population. We implement LiRa's key components and perform extensive over-the-air experiments. While LiRa inherits uplink coverage from Wi-Fi, we demonstrate that a commercial infrared uplink is subject to deep rotational fades and outages. Finally, we show that in typical WLAN scenarios, LiRa reduces response delay up to a factor of 15 and reduces throughput degradation of legacy Wi-Fi from an excessive value of 74% to less than 3% compared to transmission of VLC feedback via 802.11 without LiRa.**

I. Introduction

Visible Light Communication (VLC) is a key emerging communication technology that dual purposes LED-based lighting infrastructure for both illumination and communication. In particular, ceiling-mounted luminaries can modulate lighting in a manner unnoticeable to the human eye (i.e., flicker free) for reception at mobile devices fitted with low-cost photo diodes integrated with their casing surfaces. The proliferation of VLC-enabled luminaries promises not only support for low-rate IoT applications [21], [24] but also Gigabit rate wireless networking [3], [23]. Moreover, VLC has been demonstrated to enable higher resolution localization than radio [10], [27], [28]. Unfortunately, the wide coverage and relatively high transmit power realized by the downlink to satisfy the illumination objective is problematic to realize on the uplink: even if a mobile client is fitted with infrared LEDs, providing wide aperture long-range transmission is ill-suited to mobile devices' form factor and energy constraints [14], [15].

In this paper, we make the following contributions. First, we present the design of LiRa, a **Li**ght-**Ra**dio WLAN that fuses light and radio links on a frame-by-frame basis at the MAC layer. Unlike commercial systems [17], LiRa does not require uplink infrared transmission by the mobile client, and instead employs a radio uplink, seamlessly integrating with legacy Wi-Fi. We describe the hardware and software architecture of LiRa along with the network deployment scenario and protocol stack.

Second, a key challenge of LiRa is realizing a radio feedback path via Wi-Fi for both acknowledgements of VLC data transmissions to control ARQ and client transmission of control information such as signal strength reports, needed for luminary selection and adaptation of modulation and coding schemes. In legacy Wi-Fi, the ACK is *protected* by the same NAV that protects the ACK's associated data. Thus, a legacy Wi-Fi ACK does not separately compete for access to the medium and its transmission time is part of the duration-field that specifies the time for other stations to defer [6]. Due to LiRa's co-existence with legacy Wi-Fi, there might be an ongoing transmission in the Wi-Fi channel that prevents a client from sending feedback immediately after VLC data reception. Consequently, LiRa cannot employ the 802.11-style two way DATA-ACK handshake for downlink VLC data as such an approach would incur severe ACK delays, orders of magnitude greater than SIFS, and would consume excessive resources on the RF network.

Thus, we design AP-Spoofed Multi-client ARQ (ASMA) as a mechanism to minimize the impact of VLC control frames on legacy Wi-Fi traffic, while providing sufficient feedback for the visible light downlink. ASMA relies on three principles: First, the AP triggers VLC ARQ feedback as opposed to a traditional ACK being triggered by the reception of a certain number of bytes or frames (ACK and block ACK respectively). Moreover, VLC ARQ feedback information is opportunistically aggregated up to the time of the trigger. Second, we spoof the AP's 802.11-compliant trigger message with a sufficient NAV duration (time commanding other stations to defer) to enable *multiple* clients to transmit feedback. To avoid each of those clients independently contending, we equip the trigger message with a transmission order such that the clients can sequentially transmit during the protected NAV time. Third, the AP contends to transmit the trigger command after a configurable *feedback trigger time* which can be set to balance feedback responsiveness with airtime used in legacy Wi-Fi.

Finally, we implement LiRa and ASMA on a software defined radio [8] and perform an extensive set of experiments using a mix of LiRa components and commercial systems as baselines. By employing Wi-Fi for the uplink, LiRa's uplink is not subject to rotational outages. For comparison, we perform measurements of uplink coverage using a state-of-the-art commercial system [17] and show that its infrared uplink is subject to deep rotational fades yielding zero throughput with rotation beyond $\pm 15^\circ$ in comparison to perfect alignment. While [17] is intended for desktop rather than mobile coverage, extending uplink infrared coverage to 360° would require a bulky array of infrared LEDs. Next, we evaluate ASMA's ability to obtain channel access in a busy Wi-Fi environment and evaluate the impact of the feedback control channel on fully backlogged Wi-Fi traffic. As a baseline, we implement a feedback control channel termed Per-Client-Contention (PCC) in which each

VLC client must independently transmit its feedback via encapsulated Wi-Fi. We collect over-the-air measurements of the feedback delay and find that in a typical WLAN scenario, ASMA reduces response delay by a factor of 15 compared to PCC and reduces the impact of feedback messages on legacy Wi-Fi throughput from an excessive value of 74% for PCC to 3% for ASMA.

The remainder of this paper is organized as following. In Section II, we present LiRa's architecture. In Section III, we present LiRa's protocol design. In Section IV, we describe LiRa's implementation and the data collection. In Section V, we discuss the key results from our VLC coverage analysis and evaluation of LiRa. In Section VI, we review related work and the paper is concluded in Section VII.

II. LiRa Architecture

In this section, we present an overview of LiRa's hardware and software architecture, along with an example deployment scenario and protocol stack.

A. Hardware and Network Architecture

We design LiRa as an indoor WLAN that dual purposes luminaries for communication. A typical deployment scenario as illustrated in Fig. 1 has multiple LED lighting sources used to illuminate a room. The luminaries are typically distributed spatially solely for illumination objectives in order to avoid large shadows associated with a single point source. The LiRa AP controls these LEDs and can either group multiple sources together as a single transmission (e.g., to provide robustness for a high mobility scenario) or divide the coverage area into separate collision domains (e.g., for an auditorium scenario). The latter can be achieved using a variety of techniques including wavelength division (see [7] for a discussion). In either case, the AP transmits, but does not receive, via VLC, as the LiRa AP is not equipped with photo diodes. For the uplink, the LiRa AP receives via legacy Wi-Fi hardware and custom software as described below.

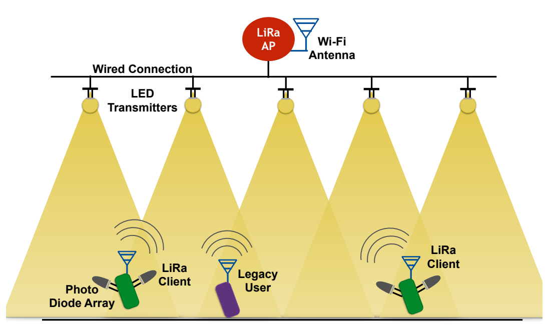

Fig. 1: Example LiRa scenario.

LiRa clients are equipped with at least one photo diode for reception of VLC signals and preferably have an array of photo diodes on multiple surfaces of the device. The photo diode array provides robustness to device orientation and minimizes the probability of blockage. For example, a device with a single photo diode that is temporarily oriented towards the floor would have a poor reception data rate, receiving only reflected signals, compared to a device that has a photo diode on each surface and can select or combine signals from multiple photo diodes. The LiRa client uses VLC for all downlink data receptions barring outage or failure and uses legacy Wi-Fi hardware and custom software for both uplink control (such as ACKs and channel state reports) and data, and for ACKs of uplink data.

Lastly, as also depicted in the figure, the LiRa AP supports legacy Wi-Fi clients which do not have VLC reception capabilities. A key component of LiRa's design is ensuring that such clients, as well as other nearby Wi-Fi networks (not shown), obtain a controlled share of airtime when interacting with a LiRa network.

Thus, the scenario and hardware architecture target to exploit the inherent coverage of downlink VLC realized by both illumination objectives and low-cost client-side photo diode reception arrays. LiRa does not attempt to realize a robust wide aperture uplink, as to do so would require a client LED transmit array along with transmit power and illumination intensity that would be significantly less than what is viable on the downlink [14], [15]. Instead, LiRa employs a radio uplink.

B. Software and Protocol Stack

LiRa provides a side-by-side light-radio protocol stack integrated via a common IEEE 802.2 interface. Consequently, as illustrated in Fig. 2, LiRa provides an abstraction of a single layer-2 hardware interface to higher layers. This approach contrasts with prior work which treats light and radio as separate networks in much the same way mobile clients have both cellular and WLAN interfaces today. The unified interface is critical to realizing both a fast VLC feedback channel and to control the impact of VLC control traffic on legacy Wi-Fi.

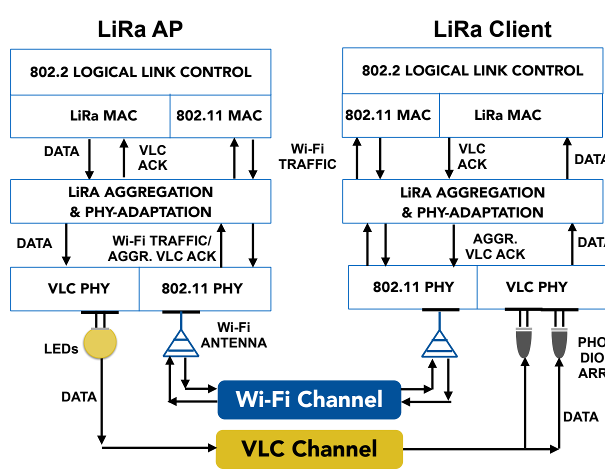

Fig. 2: LiRa's protocol stack and data flow.

From the data flow perspective, *downlink* VLC data originates from the AP only, and the primary functions of the VLC downlink MAC are scheduling, framing, and PHY adaptation. The VLC MAC is not contention based as ambient light sources not controlled by the AP are considered to be interference. The LiRa AP adapts PHY parameters such as the selection of the luminary source(s) and the modulation and coding scheme (MCS), which are impacted by client and environmental mobility as well as interference. The control feedback discussed below provides the required input for this adaptation. LiRa can employ any physical layer, including [3], [5], [23].

The LiRa client receives downlink data via an array of photo diodes. LiRa can also employ any physical layer reception mechanism that is compatible with the transmitter. Frames addressed to the client are processed up the protocol stack as illustrated in the figure.

Uplink transmissions can be divided into control frames and data frames. A data frame is transmitted via legacy Wi-Fi in the same way that legacy stations transmit. Moreover, uplink data is acknowledged by the LiRa AP using *Wi-Fi* to maintain backward compatibility and since Wi-Fi already protects the channel for the ACK that follows data. That is, there would be no advantage to using downlink VLC to ACK uplink Wi-Fi data.

Uplink *control* is quite different. Because the VLC and Wi-Fi channels are operating asynchronously, LiRa clients cannot immediately ACK a received VLC frame over Wi-Fi without risking collision or excessively disrupting ongoing traffic. Likewise, if the client contends for Wi-Fi channel access to transmit VLC ARQ feedback, the delay could be excessive and the feedback load could create heavy contention on the radio channel. Consequently, the AP controls the *time* at which ACKs and other control information are transmitted as described in Section III. Here, we describe the structure of a VLC ARQ feedback: Because client feedback is sent on-demand in response to the AP, we opportunistically aggregate ACK information up to the time that the command is received by the client. When commanded, the client sends an aggregate ACK with a bitmap representation of the frames received along with the sequence number of the first frame. This representation is similar to an 802.11 Block ACK representation [6] with the key difference that the LiRa client does not negotiate a fixed block size with the AP. Instead, this size is opportunistically set by the client at the latest possible instant, in response to the AP's command. Upon receiving the VLC ARQ feedback, the AP can then perform traditional ARQ. Likewise, we define a field in this same VLC ARQ feedback message to provide hints for the VLC transmitter to adapt its physical layer parameters. The AP can then use a combination of the loss profile (ACK bitmap) and receiver measurements such as SNR for different luminary sources to optimize downlink transmission.

Lastly, we note that with a high data rate in hundreds of Mbps for the VLC downlink and a VLC ARQ feedback delay in the order of tens of milliseconds, the buffering cost at the AP is in the order of a few megabytes. This is a modest overhead given that state-of-the-art APs are equipped with hundreds of megabytes of RAM and gigabytes of flash storage, e.g., Synology RT1900ac router.

III. A Scalable Feedback Channel for Light

In this section, we present **AP**-**S**poofed **M**ulti-client **A**RQ (ASMA), a Wi-Fi compatible feedback channel that is triggered by the AP via a spoofed NAV that protects the RF channel for a sufficient time for multiple LiRa clients to send contention-free feedback. The trigger time is managed to balance the need for timely feedback with the airtime that will otherwise be used by legacy Wi-Fi and uplink data.

A. AP Trigger

The VLC downlink transmits multiple frames in succession that can span multiple LiRa clients before reception of acknowledgement feedback to control ARQ. As illustrated in the simplified timeline of Fig. 3, VLC downlink frames are transmitted by one luminary or a group of luminaries according to policies as described in Section 2. The AP is shown sequentially transmitting variable sized frames to different LiRa clients as indicated by the depicted numbers. ASMA sets a *feedback trigger timer*: once the timer expires, the AP will aggressively attempt to access the radio channel once it becomes idle in order to transmit an ASMA trigger message.

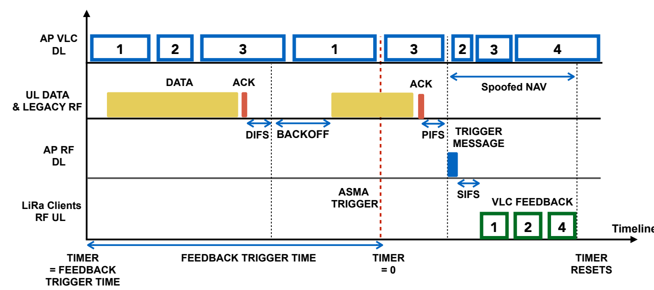

Fig. 3: Simplified LiRa timeline illustrating the combination of ASMA trigger to the AP, trigger message transmitted by the AP and uplink VLC ARQ feedback transmissions.

The expiration of the feedback trigger timer is also depicted in the figure. In the example, a Wi-Fi transmission was ongoing when the timer expired in the figure. Consequently, the AP waits until the transmission completes in order to access the channel. To minimize the channel access time, the AP employs a prioritized access strategy as follows: When a channel becomes idle, based on the 802.11 standard, each user begins backoff only after the channel becomes idle for DIFS (= SIFS + 2*SLOT) duration, where SLOT refers to the length of one backoff counter slot duration. To guarantee that the AP acquires the channel before other users start backoff, the AP sends the ASMA trigger message after the channel becomes idle for PIFS (SIFS + SLOT) time. In case other APs are in range, a less aggressive policy can be used to minimize the chance of collisions of such messages. If ARQ feedback packets from a client suffer collision, LiRa AP can re-schedule the feedback from this client.

B. Spoofed NAV for Control Channel

The objective of the ASMA trigger message is two fold: first, it protects the radio channel for a sufficient duration to enable transmission of the required feedback; second, it provides a transmission schedule for feedback such that multiple clients can send feedback messages without incurring per-client contention and collisions.

We achieve the former objective via the virtual carrier sense mechanism of Wi-Fi in which other stations defer according to the duration contained in the header's Network Allocation Vector (NAV). Namely, the LiRa AP spoofs the NAV: instead of inserting the AP's actual frame transmission duration, it advertises a NAV to enable receipt of the feedback that the AP requires from the stations. This is illustrated in the timeline which depicts the duration field indicating an interval for both the trigger message and the three feedback messages. The trigger message can be realized as a data frame or a CTS-to-self [6] which LiRa clients are programmed to recognize as a trigger message. We employ the latter approach in our implementation. Upon receiving the message, legacy stations will defer and LiRa stations will decode the trigger message.

C. Multi-Client Scheduled Feedback

The second function of the ASMA trigger is to provide a transmission schedule for client feedback messages. In particular, LiRa targets scaling the feedback process by avoiding per-client contention for each feedback message. Consequently, VLC ARQ feedback and channel state information can be efficiently communicated on the uplink in order to guide the AP's ARQ processes and PHY adaptation.

Thus, the trigger message includes an identifier and start time for each LiRa client that the AP requires feedback from. The start time, expressed in mini-slots off-set from the end of the trigger message, enables a group of clients to transmit on the radio uplink sequentially without random access or polling. The figure's timeline illustrates an example in which three clients are commanded by the trigger to transmit feedback, client 1 and 2, which have received data and will feedback both VLC ARQ feedback information and PHY updates, and client 4, which has not received data, but the AP may require other control information from, such as a PHY update to ensure that the client is matched with the best luminary.

D. Balancing LiRa Responsiveness and Control Traffic Airtime

As described above, the feedback trigger time must balance responsiveness for downlink ARQ and PHY adaptation with the airtime overhead on the Wi-Fi channel.

Wi-Fi transmissions (legacy RF and LiRa uplink data) following the 802.11 protocol occur for the entire duration of the the feedback trigger time without any use of airtime by LiRa. After the trigger expires, the LiRa AP waits at most the transmission opportunity limit for an ongoing transmission to complete. If the channel was already idle when the timer expires, the AP sends the trigger immediately.

Once the AP accesses the channel to transmit the trigger message after PIFS, the channel occupancy time includes the time to send the trigger message (denoted as T_{tm}) and the time for all LiRa clients to transmit feedback. ASMA includes a SIFS duration between the trigger message and the first LiRa client uplink transmission. This spacing is for the client scheduled first to decode the trigger message preceding its VLC ARQ feedback transmission. Thus, the total time per cycle used for feedback is at most $PIFS+T_{tm}+SIFS+NT_{fb}$, in which N is the number of LiRa stations, T_{fb} is the maximum per-station feedback time, including PHY preambles and all control fields, and PIFS and SIFS are the same Wi-Fi standard values. Thus, denoting T_{ftt} as the feedback trigger time, LiRa utilizes no more than a fraction

$$1-\frac{T_{ftt}}{T_{ftt}+PIFS+T_{tm}+SIFS+NT_{fb}} \qquad (1)$$

of Wi-Fi airtime.

While Equation (1) favors a large feedback trigger timer in order to amortize overhead and minimize the impact on legacy and uplink Wi-Fi, a smaller value is favored for ARQ and PHY responsiveness. The maximum delay for feedback is the sum of the trigger time, the maximum Wi-Fi transmission time $\text{TXOP}_{\max}$, and the maximum feedback time as given above, i.e.,

$$T_{ftt}+\text{TXOP}_{\max}+PIFS+T_{tm}+SIFS+NT_{fb}. \qquad (2)$$

Thus, the dominant term in each case is the feedback trigger time with an inverse relationship for fraction of airtime and a linear impact on response delay. We evaluate policies for setting the trigger time in practice in our experimental evaluation.

IV. LiRa Implementation

In this section, we describe our implementation of the key components of LiRa, including a custom VLC platform to collect measurements in a typical indoor environment, and a software-defined radio implementation of LiRa's RF components.

A. VLC Downlink Implementation

VLC AP transmitter. As described in Section 2, LiRa can employ any VLC physical layer. Here, we repurpose Philips Hue smart-LED lightbulbs [16] that are capable of changing hues and light intensities as transmitters. While the Philips' hardware and software interfaces do not allow high frequency modulation, the platform enables study of a large set of LiRa performance factors as described below. Most critically, the achievable data rate on a VLC link is a function of the received signal strength at the receiver's photodiode. This signal strength depends on the luminary's transmit power, the distance between luminary and photo diode, the incidence angle from the light source, and the irradiation angle at the photodiode. The Philips system along with most other commercial LED luminaries are equipped with diffusers that aid to realizing uniform propagation, thereby minimizing the impact of incidence angle.

VLC client receiver. The LiRa VLC receiver employs Adafruit TSL2591 high dynamic range digital light sensors. We mount the light sensor on top of an Arduino platform that handles serial communication with a computer tasked with data processing. We place the receiver on a transparent holder to reduce undesirable blockage from nearby. This holder is designed to have two concentric acrylic layers that can be rotated 360° relative to each other. This structure is combined with a motion controller [2] that consists of a slider along which the receiver can be linearly moved over 50 cm, and motors that can be controlled via MATLAB to provide desired movements and rotations for the receiver.

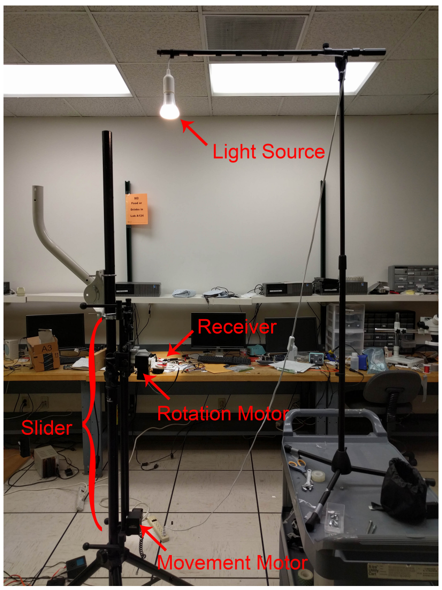

Fig. 4: Smart LED-based experiment setup.

Experimental setup. Fig. 4 illustrates our baseline experimental setup. For all measurements, we hang the smart-LED lightbulb overhead with the help of a stand with an adjustable height. We collect light intensity measurements at locations corresponding to different distances between the LED bulb and receivers spanning over 150 cm. At every location, the receiver's sensor is rotated in steps of 10°, with the sensor's rotation corresponding to the varying irradiance angle that determines the light intensity received. We collect 168 samples of light intensity for every location and every sensor rotation angle. Using this measurement database and 802.15.7 standard MCS table, we compute the MCS the AP would select for every location and orientation of the sensor.

B. Radio Link Implementation

Implementation. We utilize WARP v3 [8] for the radio link measurements. The WARP board exchanges data with a computer via the Ethernet interface and can perform real-time wireless transmission with its WARPNet module. We use the 802.11 reference Design for WARP v3 as the MAC layer design, which is 802.11g compatible. We utilize multiple WARP boards, including 1 AP and different combinations of VLC users and legacy users.

With ASMA, the AP sets the feedback trigger timer after its last reception of feedback and attempts to obtain channel access after the timer expires. During this period of time, the VLC traffic of LiRa has no impact on the Wi-Fi channel. Once the AP obtains channel access for the trigger, VLC ARQ feedback transmissions are contention free. In our implementation, the AP transmits the trigger message over the air along with the spoofed NAV field. In order to realize the feedback trigger message, the LiRa AP targets to obtain channel access as soon as possible after the channel becomes idle upon expiration of the trigger timer. As specified by ASMA, we replace the DIFS value of the AP with PIFS and set the contention window size to 0 whenever a trigger message should be transmitted.

The trigger frame's payload must identify the association ID and time for each of the requested feedback messages. Consequently, we set the trigger frame payload size accordingly.

We implement fully backlogged legacy users, which represents the worst case for the AP competing to win contention for the trigger message. We focus on the VLC ACK feedback analysis and therefore do not transmit Wi-Fi uplink data traffic for LiRa clients and instead consider the impact of such traffic to be equivalent to the traffic of other legacy users.

Experimental Setup. To evaluate the impact of feedback trigger time, LiRa client size, and legacy user traffic characteristics, we collect over-the-air measurements using the WARP-based LiRa implementation. We couple the above measurements with computations that utilize the traces from our VLC downlink measurements as input. We generate 1,000 client locations and assign each client a location and orientation randomly following a uniform distribution of the locations and orientations used in the smart LED bulb measurement study described above. For each location, we compute the per-client MCS that would be selected by the AP for the downlink transmission given the measured signal strength. Then, for different feedback trigger times, we compute the trigger frame payload size for transmission via WARP.

Each experiment run consists of a wall-clock time of 10 seconds in which thousands of legacy user data packets and LiRa feedback messages are transmitted. For analyzing Wi-Fi traffic's impact on LiRa's performance, we perform 20 independent reruns spanning several hours given a setting of LiRa trigger time, Wi-Fi operating channel and legacy user traffic flows. We transmit all legacy user data packets via IEEE 802.11g compliant operation and with all frames having length of 1024 bytes with variable MCS unless stated otherwise. The AP has fully backlogged data for the LiRa clients. First, we collect measurements for different feedback trigger times and different trigger frame message sizes corresponding to a unique pairing of feedback trigger time and LiRa client size. Second, we collect measurements for different uplink MCS of the legacy users and the operating Wi-Fi channels to analyze the impact of legacy user traffic on LiRa's performance. Third, for baseline comparisons, we implement other feedback strategies as described in the next section.

C. Configuration

Unless otherwise noted, all results use the following configuration: First, the ASMA trigger message and the LiRa client's VLC ARQ feedback is transmitted at the base rate of 6 Mbps and the VLC downlink MPDU (MAC Protocol Data Unit) is 1 kB.

Second, the number of ACKs opportunistically aggregated into a VLC ARQ feedback transmission depends on the VLC data rate used by the AP to serve the client, as a higher rate yields more ARQ feedback. The VLC downlink rate is dependent on both the PHY architecture of the VLC link as well as environment-dependent factors such as distance. Unless otherwise noted, we map measured signal strength to MCS as specified by 802.15.7, and incorporate all MAC and PHY layer parameters from 802.11 and 802.15.7 standards for the Wi-Fi and VLC links respectively.

Finally, the AP has fully backlogged traffic for downlink transmission to each LiRa client and the legacy users have backlogged traffic for the AP. For scenarios with multiple LiRa clients in the network, the AP conducts the VLC transmission using round-robin scheduling of frames. More sophisticated scheduling [13] and rate adaptation [29] can be applied to improve LiRa performance.

V. Experimental Evaluation

In this section, we evaluate the performance of LiRa using the above implementation platform. We study VLC coverage, response delay and VLC feedback's impact on legacy Wi-Fi traffic for a broad class of scenarios and configurations.

A. Uplink Coverage: Limits of Infrared

Because LiRa employs a radio uplink, here we evaluate limits of employing *infrared* (or similarly, visible light) as an uplink. In particular, we utilize a commercial system, pureLiFi [17], that, like LiRa, provides downlink coverage via LED light fixtures connected to an AP. However, in contrast to LiRa, pureLiFi's uplink uses infrared instead of radio. The pureLiFi client does not employ an infrared LED transmitter on each of its surfaces; indeed, we expect that due to the relative bulk of LED transmitters (similar to the size of a phone's flash), it would be infeasible to place one on each of the mobile client's surfaces, even in future designs.

PureLiFi enables separate locations for the AP's LED bulb acting as the luminary and downlink transmitter and the AP's photo-diode based infrared receiver. For our purposes of evaluating coverage, we place these two AP components adjacent to each other and position the client 130 cm below as illustrated in Figure 5. We study uplink coverage and the impact of the client's infrared transmitter orientation on uplink throughput and vary the orientation of the client about its edge axis The orientation is changed in steps of 5° with 0° representing perfect alignment (i.e., the infrared client transmitter is perfectly aligned with the AP's ceiling mounted photo diode receiver). We perform the rotation in both clockwise and counter-clockwise directions. For each orientation angle, we conduct 10 independent runs and use *iperf* to measure the

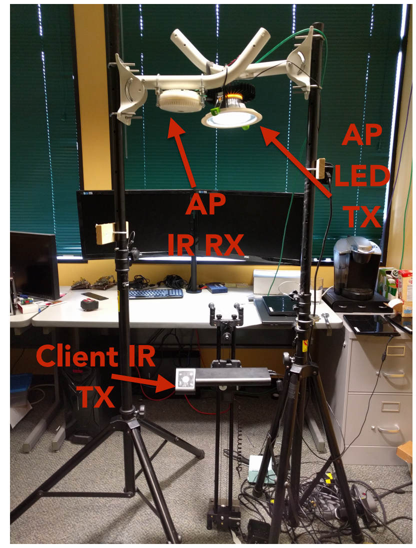

Fig. 5: pureLiFi setup where IR stands for infrared.

data transferred in an interval of approximately 10 seconds under TCP protocol with default TCP window size of 85.3 kB.

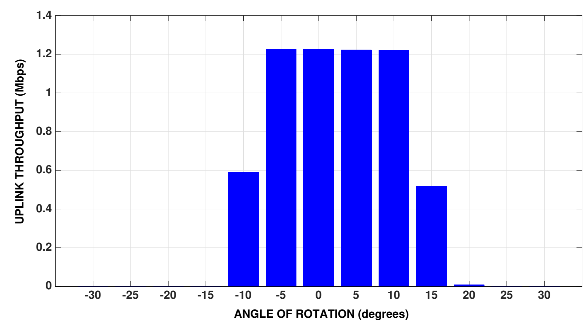

Fig. 6: Uplink throughput of pureLiFi versus orientation angle.

The results are depicted in Figure 6, which shows uplink throughput vs. orientation angle. Negative and positive angles indicate clockwise and counter-clockwise rotations respectively. First, the throughput is close to the peak value at the best alignment and $\pm 5°$. Second, with increased rotation of 10°, the throughput remains at 1.2 Mbps in the counter-clockwise direction but reduces to 0.6 Mbps in the clockwise direction. This asymmetry arises due to the increased distance between the transmitter and receiver in the case of clockwise rotation as the rotation axis is the edge of the client's device. *Finding: At 20° of rotation, we observe negligible throughput of 8.4 kbps, beyond which the link can no longer be established, yielding zero throughput.* In contrast, LiRa targets mobile clients and consequently employs radio as an uplink medium, which is not subject to such deep rotational fades and outages.

B. Response Delay Evaluation

Here, we analyze the impact of feedback trigger time, the number of LiRa clients, and interference from Wi-Fi traffic on response delay. The response delay is the duration between when the AP transmits a VLC frame to a particular client and when the AP receives ARQ feedback for that particular frame. The *mean* response delay is an average over all frames and clients, as well as over time with multiple trigger messages. Likewise, the feedback trigger time is the duration of the timer that the AP uses before contending for access to the wireless channel to send the trigger message, as described in Section III-A.

In ASMA, the AP wins channel contention by transmitting the trigger message a duration PIFS after the channel becomes idle. However, the time between ASMA triggering the AP to enqueue the trigger message for transmission and the AP transmitting the trigger message over-the-air depends on the Wi-Fi traffic characteristics which includes uplink and downlink transmissions from legacy stations as well as uplink data from LiRa clients. Moreover, even when a controlled experiment is run with a fixed number of legacy users, there might be interference from legacy users associated with other APs nearby. For this purpose, we analyze the impact of Wi-Fi traffic on ASMA's response delay in two dimensions: number of Wi-Fi traffic flows and the Wi-Fi operating channel. In our experiments, there are either one or three Wi-Fi traffic flows in the network and we set their uplink PHY data rate to 18 Mbps. We fix the feedback trigger time to 5 ms and the number of LiRa clients to be a high value of 10. We analyze the impact of co-channel interference by conducting the experiments in three different Wi-Fi channels of operation: channels 1 and 14 in the 2.4 GHz band and channel 48 in the 5 GHz band.

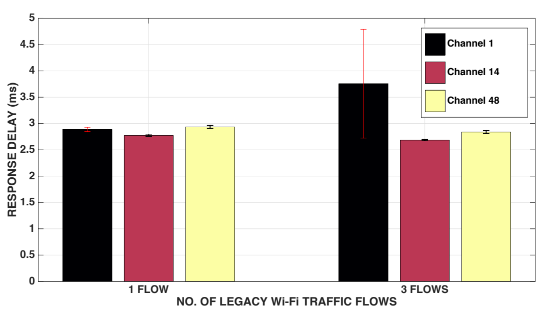

Fig. 7: Wi-Fi traffic impact on LiRa's response delay: 1 flow and 3 flows.

Fig. 7 depicts mean response delay as a function of the combination of the number of traffic flows and the Wi-Fi operating channel. First, independent of the number of traffic flows and the Wi-Fi channel, the mean response delay is consistently lower than the corresponding trigger time value. This is because frames transmitted in the latter part of the trigger duration typically have a response delay that is less than the trigger time itself (cf. Fig. 3). Second, there are differences in response delay across different channels due to uncontrolled Wi-Fi transmissions occurring in the range of the AP. The channels 14 and 48 have negligible variance in response delay because they are not used for commerical Wi-Fi operation in the experiment coverage area. Finally, response delay increases with the number of flows in Channel 1. This is due to the increased probability of the channel being occupied by legacy user transmissions when the trigger timer expires. In our experiments, we observe behavior similar to this figure for varying sets of LiRa client sizes and trigger times spanning from 1 ms (frequent feedback) to 10 ms (less frequent feedback). *Finding: The average response delay is lower than LiRa's feedback trigger time independent of LiRa client size and legacy Wi-Fi traffic.*

C. Feedback with Per-Client Contention

To analyze the gains provided by ASMA's strategies of spoofed NAV and multi-client scheduled feedback, we compare ASMA's performance with an alternative strategy employing client-driven feedback.

We define *Per-client Contention (PCC)* as a feedback mechanism in which each VLC client *independently* contends via 802.11 to transmit feedback, such that VLC feedback is treated as an encapsulated Wi-Fi data frame. This approach contrasts with ASMA in that because the AP has not provided a spoofed NAV to allow feedback, the PCC clients providing feedback must contend independently. Consequently, the total feedback contention on the radio channel is once per client vs. once per trigger time. Nonetheless, we do not require PCC clients to contend for each downlink VLC *frame*. Instead, a PCC client begins contention as soon as it has feedback to send to minimize the response delay. Also, each PCC client opportunistically aggregates ARQ feedback up to the time that it obtains channel access for sending the feedback message. Thus, similar to LiRa, PCC uses a bitmap representation with opportunistic aggregation of feedback information up to the time that the client transmits.

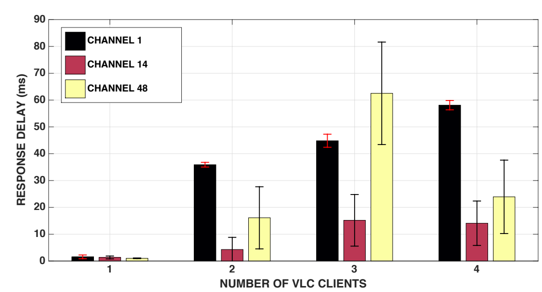

Fig. 8: Response Delay performance of Per-Client Contention in different Wi-Fi channels.

We configure a single legacy user with fully backlogged traffic for the AP and a varying number of VLC clients. Fig. 8 depicts the average response delay of PCC vs. the number of VLC clients in different Wi-Fi channels. First, when there is a single client in the network, the response delay is less than 10 ms in all channels comparable to LiRa's response delay. When the VLC client size increases to 2, the response delay in channel 1 increases to 35 msec. This is due to the uncontrollable, ongoing transmissions in this channel during this experiment run. In contrast, the response delay in channels

14 is still below 10 msec. In this channel, not only does the client have negligible interference from other transmissions but also there is a high probability of winning contention as there is just a single legacy user associated to its AP. Second, with an increase in the number of clients, PCC's response delay in all channels increases highlighting the strong impact of contention on the response delay performance of PCC. This increased delay results from the airtime lost in collisions among the PCC feedback and with the uplink data packets of legacy user. Third, increasing from three to four clients results in a decrease in mean response delay for channel 48 and a negligible change in channel 14. This is because although there is increased probability of collisions, there is also increased probability of a VLC client (vs. the legacy user) winning the contention for VLC ARQ feedback transmission. In contrast, channel 1 being used for commercial operation has the additional effect of co-channel interference resulting in an increase in response delay.

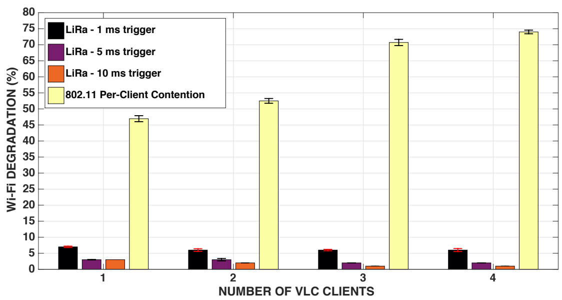

Fig. 9: Wi-Fi throughput degradation comparison between LiRa and Per-Client Contention model.

Finally, we compare PCC to ASMA and Fig. 9 depicts measured Wi-Fi throughput degradation vs. the number of VLC clients. First, as discussed in Section III, for LiRa, legacy Wi-Fi throughput degradation decreases inversely proportional to feedback trigger time. This is because the airtime required for ARQ feedback messages increases at a much slower rate than the increase in trigger time, as only one ACK bit is added for every data packet received in the downlink. For example, with a data rate of 10 Mbps in the VLC downlink, a 1kB frame requires 819.2 μsec to transmit. Corresponding to this frame, a single ACK bit added would take 0.17 μsec of the Wi-Fi airtime. Second, for LiRa, Wi-Fi throughput degradation has higher variance for short trigger times. This is because with short trigger times, the client feedback time is a significant factor (as high as 92% for 1 msec triggers) in the airtime utilized by LiRa in the Wi-Fi channel and this feedback scales linearly with the number of clients as defined in Equation (1). Third, the results indicate that PCC consistently has over 7 times the degradation of LiRa, independent of the number of VLC clients and LiRa trigger time. When there is a single VLC client, the PCC client attempts to win the channel immediately after receiving the first packet since the last PCC feedback transmission. This leads to 47% degradation of legacy Wi-Fi throughput. As the VLC client size increases, the increased occupancy by the VLC clients for their independent PCC feedback frames and the additional time lost due to collisions leads to 74% degradation in Wi-Fi throughput. In contrast, LiRa has a maximum degradation of 7% when the trigger time is as low as 1 ms. Hence, in contrast to PCC, ASMA provides a responsive and scalable feedback mechanism.

Finding: With feedback trigger time of 5 msec, LiRa can reduce the response delay by a factor of 15 and reduce the legacy Wi-Fi througput degradation to 3% from an excessive value of 74%.

VI. Related Work

VLC Only Networks. A software-defined VLC system including a bi-directional VLC link with On-Off keying modulation is presented in [24]. Likewise, Li-Flame is a commercial VLC system with 10 Mbps downlink up to 3 meters and 10 Mbps uplink via infrared [17]. Moreover, significant progress has been made towards advancing the physical layer of VLC communication for non-laser-based incoherent transmission. For example, the fastest VLC link as of this writing is a custom 3.25 Gbps system based on Single Carrier Frequency Domain Equalization utilizing an RGB LED [25]. Today's VLC standards such as IEEE 802.15.7 [5] also employ visible light for both the uplink and downlink.

In addition to WLANs, low-power devices with kbps-scale data rate capabilities have been designed for sensor networks and Internet of Things applications. Examples include a novel PHY and VLC MAC layer for energy efficient LED-to-LED communication [21], duplex, battery-free communication using a retro-reflector fabric that backscatters light [9], and OpenVLC, an open source software-based VLC research platform based on BeagleBone [24].

In contrast, LiRa overcomes inherent limitations of infrared or visible light LED-based communication applied to *uplink* WLAN access. Namely, as we experimentally demonstrated in Section V, while the illumination objective of the downlink ensures that the LED transmitters of the AP have wide aperture, large field of view, and high transmit power, LEDs on the client device have none of these benefits [14]. Namely, the limited size, power, and aperture of the mobile client's LED transmitter can severely constrain field of view, thereby limiting data rate or even breaking the uplink. Consequently, mobility and rotation of user devices might lead to significant throughput reductions or blockages.

Integrated VLC-RF Networks. VLC and RF have been jointly used in prior work: VLC was proposed as an additional directional channel to offload RF broadcast traffic in congested networks [18]; load balancing between VLC and RF interfaces was optimized in [12]; horizontal and vertical handover mechanisms between VLC and RF networks were designed in [1]; VLC transmission order was proposed to mirror Wi-Fi transmission order, as the latter will have already resolved contention [4]; routing and address spoofing at the IP layer was demonstrated to integrate VLC and Wi-Fi at layer three [22]. In contrast, LiRa is the first system to integrate VLC and RF at the MAC layer, and hence is the first system to provide a virtual feedback channel for VLC via Wi-Fi. Nonetheless, the aforementioned works are complementary to LiRa and can be used to enhance LiRa at other layers.

Centralized Scheduling. In existing centralized scheduling protocols such as 802.11 point coordination function (PCF), the AP polls a single client for data and not the ACK as the ACKs are reserved by the 802.11 mechanism. Also, the AP doesn't have knowledge of clients' backlogged traffic and therefore cannot reserve the channel for fully backlogged traffic from multiple clients. In contrast, in ASMA, the LiRa AP is able to trigger and reserve the channel for the complete duration of fully backlogged VLC feedback transmissions from multiple clients utilizing the downlink VLC scheduling information.

VLC Services and Devices. Lastly, there is an emerging body of research on employing VLC for sensing or localization and employing cameras as receivers. For example, a VLC module was designed to locate a user's finger within a workspace with one-centimeter precision [26]. Likewise, a VLC sensing system can reconstruct the 3D human skeleton postures from 2D shadow information [11]. Further, cameras were demonstrated as receivers in [19], [20]. With the addition of photo diode receivers to increase data rate, LiRa can be deployed in parallel with such applications to jointly communicate while also providing such services.

VII. Conclusion

We presented the design and implementation of LiRa, a WLAN that fuses simplex VLC downlink and bi-directional Wi-Fi on a frame-by-frame basis at the MAC layer. In order for LiRa clients to transmit VLC ARQ feedback via Wi-Fi without excessive contention-based delays, we presented ASMA as a scalable VLC feedback channel over RF. Using over-the-air measurements, we demonstrated limitations of uplink coverage using infrared/VLC. Moreover, we showed that compared to feedback using 802.11-based per-client contention, LiRa's spoofed NAV and multi-client scheduled feedback reduces response delay by a factor of 15 and reduces degradation of Wi-Fi throughput to 3% from 74%.

VIII. Acknowledgements

This research was supported by Cisco, Intel, the Keck Foundation, and by NSF grants CNS-1642929, CNS-1514285, and CNS-1444056.

References

[1] X. Bao, X. Zhu, T. Song, and Y. Ou. Protocol Design and Capacity Analysis in Hybrid Network of Visible Light Communication and OFDMA Systems. *IEEE Transactions on Vehicular Technology*, 63(4):1770–1778, 2014.

[2] Cinetics. Axis360 Pro modular motion control system.

[3] G. Cossu, A. M. Khalid, P. Choudhury, R. Corsini, and E. Ciaramella. 3.4 Gbit/s Visible Optical Wireless Transmission Based on RGB LED. *OSA Optical Express*, 20(26):B501–B506, Dec 2012.

[4] W. Guo, Q. Li, H. Yu, and J. Liu. A parallel transmission MAC protocol in hybrid VLC-RF network. *Journal of Communications*, 10(1), 2015.

[5] IEEE Computer Society. IEEE Standard for Local and Metropolitan Area Networks - Part 15.7: Short-Range Wireless Optical Communication Using Visible Light. September 2011.

[6] IEEE Computer Society. IEEE Draft Standard for Local and Metropolitan Area Networks - Specific Requirements - Part 11: Wireless LAN Medium Access Control (MAC) and Physical Layer (PHY). June 2012.

[7] T. A. Khan, M. Tahir, and A. Usman. Visible light communication using wavelength division multiplexing for smart spaces. In *Proc. of IEEE Consumer Communications and Networking Conference (CCNC)*, 2012.

[8] A. Khattab, J. Camp, C. Hunter, P. Murphy, A. Sabharwal, and E. Knightly. WARP: A Flexible Platform for clean-slate Wireless Medium Access Protocol Design. *ACM SIGMOBILE Mobile Computing and Communications Review*, 12(1):56–58, 2008.

[9] J. Li, A. Liu, G. Shen, L. Li, C. Sun, and F. Zhao. Retro-VLC: Enabling battery-free duplex visible light communication for mobile and iot applications. In *Proc. of ACM HotMobile*, 2015.

[10] L. Li, P. Hu, C. Peng, G. Shen, and F. Zhao. Epsilon: A Visible Light Based Positioning System. In *Proc. of USENIX NSDI*, 2014.

[11] T. Li, C. An, Z. Tian, A. T. Campbell, and X. Zhou. Human Sensing Using Visible Light Communication. In *Proc. of ACM MobiCom*, 2015.

[12] X. Li, R. Zhang, and L. Hanzo. Cooperative Load Balancing in Hybrid Visible Light Communications and WiFi. *IEEE Transactions on Communications*, 63(4):1319–1329, 2015.

[13] X. Li, R. Zhang, J. Wang, and L. Hanzo. Cell-Centric and User-Centric Multi-User Scheduling in Visible Light Communication aided networks. In *Proc. of IEEE ICC*, 2015.

[14] M. Noshad and M. Brandt-Pearce. Can visible light communications provide Gb/s service? *arXiv*, 2013.

[15] P. H. Pathak, X. Feng, P. Hu, and P. Mohapatra. Visible light communication, networking, and sensing: A survey, potential and challenges. *IEEE Communications Surveys & Tutorials*, 17(4):2047–2077, 2015.

[16] Philips. Hue White and Color bulb kit.

[17] pureLiFi. Li-Flame high-speed wireless network solution using VLC.

[18] M.B. Rahaim, A.M. Vegni, and T.D.C. Little. A Hybrid Radio Frequency and Broadcast Visible Light Communication System. In *Proc. of IEEE GLOBECOM Workshop on Optical Wireless Communications*, 2011.

[19] R. D. Roberts. A MIMO protocol for camera communications (CamCom) using undersampled frequency shift ON-OFF keying (UFSOOK). In *Proc. of IEEE Globecom Workshop on Optical Wireless Communications*, 2013.

[20] S. Schmid, L. Arquint, and T. R. Gross. Using Smartphones as Continuous Receivers in a Visible Light Communication System. In *Proc. of 3rd ACM Workshop on Visible Light Communications Systems*, 2016.

[21] S. Schmid, J. Ziegler, G. Corbellini, T. R. Gross, and S. Mangold. Using consumer LED light bulbs for low-cost visible light communication systems. In *Proc. of the 1st ACM MobiCom workshop on Visible Light Communication Systems*, 2014.

[22] S. Shao, A. Khreishah, M. Ayyash, M. B. Rahaim, H. Elgala, V. Jungnickel, D. Schulz, T. D. C. Little, J. Hilt, and R. Freund. Design and Analysis of a Visible-Light-Communication Enhanced WiFi system. *OSA Journal of Optical Communications and Networking*, 7(10):960–973, 2015.

[23] D. Tsonev, S. Videv, and H. Haas. Towards a 100 Gb/s visible light wireless access network. *OSA Optics Express*, 23(2):1627–1637, 2015.

[24] Q. Wang, D. Giustiniano, and D. Puccinelli. OpenVLC: Software-Defined Visible Light Embedded Networks. In *Proc. of the 1st ACM Workshop on Visible Light Communication Systems*, 2014.

[25] Y. Wang, R. Li, Y. Wang, and Z. Zhang. 3.25-Gbps visible light communication system based on single carrier frequency domain equalization utilizing an RGB LED. In *Proc. of OSA Optical Fiber Communication Conference*, 2014.

[26] C. Zhang, J. Tabor, J. Zhang, and X. Zhang. Extending Mobile Interaction Through Near-Field Visible Light Sensing. In *Proc. of ACM MobiCom*, 2015.

[27] C. Zhang and X. Zhang. LiTell: Indoor Localization Using Unmodified Light Fixtures: Demo. In *Proc. of ACM MobiCom*, 2016.

[28] J. Zhang, C. Zhang, X. Zhang, and S. Banerjee. Towards a Visible Light Network Architecture for Continuous Communication and Localization. In *Proc. of the 3rd ACM Workshop on Visible Light Communication Systems*, 2016.

[29] J. Zhang, X. Zhang, and G. Wu. Dancing with light: Predictive in-frame rate selection for visible light networks. In *Proc. of IEEE INFOCOM*, 2015.