Design and Implementation of a Hybrid RF-VLC System with Bandwidth Aggregation

Zhouchi Li,* Sihua Shao,* Abdallah Khreishah,* Moussa Ayyash,† Iman Abdalla,‡ Hany Elgala,§
Michael Rahaim,¶ and Thomas Little‡
*Department of Electrical and Computer Engineering, New Jersey Institute of Technology, Newark, NJ, USA
†Department of Information Studies, Chicago State University, Chicago, IL, USA
‡Department of Electrical and Computer Engineering, Boston University, Boston, MA, USA
§Computer Engineering Department State University of New York at Albany Albany, NY, USA
¶Department of Electrical and Computer Engineering, University of Massachusetts Boston, Boston, MA, USA
*{zl445, ss2536, abdallah}@njit.edu, †mayyash@csu.edu, ‡{imanha, tdcl}@bu.edu, §helgala@albany.edu,
¶Michael.Rahaim@umb.edu

***Abstract*—Visible light communication (VLC) has the potential to add significant capacity to short range wireless access technology by piggybacking data on light from overhead luminaires. However, an uplink is required to complete such a network, which introduces new issues. In this paper, we propose and implement a practical hybrid WiFi-VLC system that does not require a separate VLC uplink but rather aggregates WiFi and VLC downlinks and shares the WiFi uplink. Aggregated downlink bandwidth of the hybrid system is achieved by using a Linux bonding driver and media access control (MAC) address redirection. The throughput of the system is tested and compared with WiFi-only (one WiFi downlink) and asymmetric (one VLC downlink) systems under a congested WiFi environment. The evaluation results show that our system achieves aggregated downlink bandwidth that is approximately the summation of the downlink capacities of the WiFi-only and asymmetric systems. The study of the round-trip time (RTT) demonstrates the tradeoff between bandwidth utilization and latency that can be used in the design of load-balancing algorithms. Finally, the deployed system demonstrates feasibility in typical indoor space room dimensions.**

***Keywords*—Hybrid system, WiFi, visible light communication, bandwidth aggregation.**

I. INTRODUCTION

The demand for wireless network access shows no end. New applications driven by the development of Internet of Things (IoT) devices will continue to put new demand on wireless access networks [1]. However, the overuse of the radio frequency (RF) spectrum is demonstrating current and future limitations. In indoor environments, where mobile devices densely populate, WiFi users often suffer from severe contention and interference. To alleviate this problem, new approaches such as visible light communication (VLC) can be used.

The use of the optical spectrum, including VLC, is considered a reliable way to overcome the crowded RF spectrum in indoor environments. The LED-based VLC technology holds some prominent advantages, including high area spectral efficiency, high energy efficiency, high security, and dual-use nature [2]. Its easily sequestered propagation property makes ultra-dense VLC deployment possible [3] and can be used to accommodate the requirement of indoor wireless access. If VLC and WiFi are aggregated in an indoor environment; high throughput, robust transmission, and seamless coverage can be achieved. An aggregated WiFi-VLC system which combines a duplex WiFi link and a duplex VLC link has been previously demonstrated in reference [4]. However, when a duplex VLC channel is implemented, the uplink introduces problems. First, mobile devices are required to install VLC transmitters (uplinks) to send data, which increases their costs, complexity, and power consumption [5]. Second, alignment of an optical uplink is required from each mobile device transmitter to a receiving AP which requires transmitter beam angle control and manipulation of the receiver field of view (FOV). Third, the VLC uplink channel can generate uncomfortable glare to humans if it is in the visible bands [6]. These problems can be overcome by using the existing RF (WiFi) channel as an uplink to the AP.

When an RF uplink is used instead of the VLC uplink in the hybrid RF-VLC system, downlink traffic scheduling becomes an issue. This is due to both the VLC and the WiFi channel traffic sharing the same uplink. Packets on the shared uplink must be interpeted so as to generate downlink traffic on the correct medium–VLC downlink or WiFi downlink.

In this paper, we propose and implement a practical hybrid WiFi-VLC system comprised of two duplex links. One is a duplex WiFi link and another is a VLC-based asymmetric link paired with a WiFi uplink. The WiFi uplink of VLC and the duplex WiFi link transmit data through the same channel. This WiFi uplink makes the system effective because the VLC uplink is avoided. The system achieves aggregated bandwidth on the data link layer by using the Linux bonding driver [7] under mode 6 on the client. Under this mode the Linux bonding driver builds a logical interface on the client. All network interfaces on the client are controlled by the logical interface and they are not required to have Internet Protocol (IP) addresses. The logical interface has the IP address and uses the media access control (MAC) address of one of the network interfaces. The logical interface schedules the traffic among the network interfaces and all devices that

978-1-5386-2070-0/18/$31.00 ©2018 IEEE

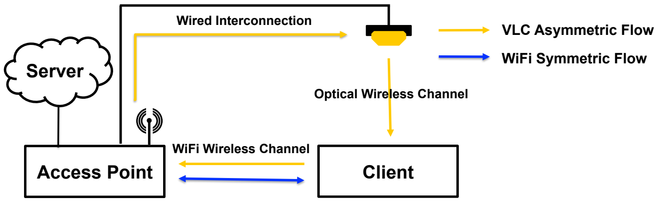

Fig. 1. Proposed Hybrid WiFi-VLC System

communicate with the client only have the knowledge of the logical interface. Data link layer bandwidth aggregation provides higher utilization of the aggregated capacity [8]. Fig. 1 shows the basic configuration of this hybrid system. The main contributions of this paper are summarized as follows:

- A hybrid WiFi-VLC system is designed by utilizing the uplink of WiFi to avoid the undesirable characteristics of a VLC uplink.
- The hybrid WiFi-VLC system is implemented and demonstrated, showing the benefits of aggregation.
- The throughput and round-trip time (RTT) of the hybrid WiFi-VLC system is evaluated under different WiFi contention scenarios and VLC operating distances, and benchmarked against a conventional WiFi system and a hybrid WiFi-VLC system. The tradeoff between the bandwidth utilization and latency in algorithm design of the hybrid system is analyzed by considering the RTT. The impact of the VLC operating distance on the system throughput is also studied to show the feasibility of the hybrid system in an indoor environment.

The remainder of this paper is organized as follows. Section II reviews some important related works on hybrid WiFi and VLC systems. Section III briefly introduces the architectures of the existing VLC asymmetric system and aggregated system implemented in references [9] [4], as well as our hybrid system. Section IV provides analysis and experimental results demonstrating the benefit of the proposed hybrid system. Finally, Section V concludes the paper.

II. RELATED WORK

In reference [9], an asymmetric hybrid WiFi-VLC system is designed and implemented in order to leverage the downlink capacity of VLC. In this system, request packets sent by the client are transmitted through the WiFi channel. The reply packets sent by the server are received via the VLC channel. However, under this model, the WiFi channel is only utilized as the uplink of the client, suggesting that more performance could be realized if the downlink capacity were aggregated with the VLC capacity.

In reference [6], an indoor hybrid system that integrates WiFi and VLC luminaries has been presented and simulated. In this system, broadcast VLC channels are utilized to supplement RF communications. A handover mechanism between WiFi and VLC is designed to dynamically distribute resources and optimize system throughput. However, the analysis is based on the assumption of a reliable WiFi uplink and has not been validated in a real implementation.

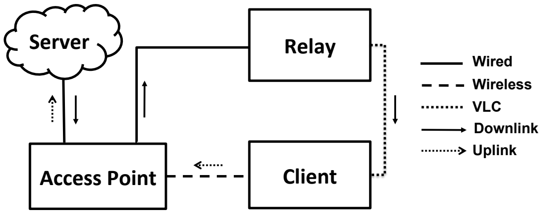

Fig. 2. Asymmetric System Architecture

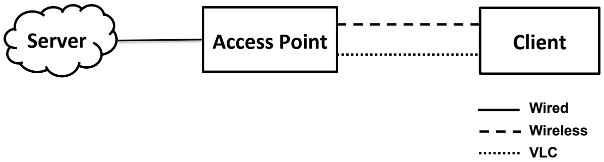

Fig. 3. Aggregated System Architecture

In reference [4], an aggregated system with a duplex WiFi and VLC link is implemented and analyzed. The system can achieve high throughput performance by not only providing aggregated bandwidth, but also load balancing between different interfaces. Since the VLC uplink channel can produce an unpleasant irradiance glare from the user device, it is unlikely that this kind of configuration will be adopted in practice.

In this paper, we expand upon our previous publications in order to show a practical WiFi-VLC hybrid system that requires the same number of WiFi and VLC APs as found in references [9] and [4]. However, the proposed system achieves aggregated bandwidth without requiring a VLC uplink.

III. SYSTEM MODEL

In this section, we briefly introduce the architectures of two WiFi-VLC hybrid systems from which we develop our hybrid system. Subsequently, we present our practical hybrid system (Fig. 4) that integrates two duplex links. The challenges of the hybrid system are described below. The presented architecture of the hybrid system, which might not be the optimal design, is primarily selected for proof-of-concept demonstration of the concepts discussed in this paper.

A. Asymmetric System

Fig. 2 shows the architecture of the asymmetric system. Here, the uplink and downlink flows are transmitted through different channels. For the uplink, data are transmitted from the client through WiFi channel. For the downlink, data are received by the relay first and then transmitted to the client via a VLC channel.

B. Aggregated System

The architecture of the aggregated system is shown in Fig. 3. There are two duplex links between the client and the AP. One uses the WiFi channel and the other uses the VLC channel. Aggregated bandwidth is achieved through this system.

C. *Challenges*

The primary challenges of the hybrid system implementation are as follows:

1) **The reply data from the server must be received by the network interface card (NIC) which sends the corresponding request packets.** Every time a Transmission Control Protocol (TCP) connection is built, the client initiates a three-way handshaking with the server. The client generates an SYN segment in the application layer and the segment is encapsulated within an IP header before it is sent through the NIC. The client listens to the socket with the TCP port and IP address which are used to generate and encapsulate the SYN segment. If the reply packets from the server are received by a different NIC with different socket information, the application which builds the TCP connection will not process the reply packets [4]. As shown in Fig. 4, there are 3 NICs on the client. In the asymmetric VLC link, NIC W-1 is used to transmit data and NIC B-2 is used to receive data.

The client can be forced to use the same NIC to transmit and receive data by deleting all information of other NICs in the routing table of the client. Thus, the client will check its routing table and use the rest of the NIC in the routing table to send and receive data. However, since the Linux bonding driver is run under mode 6 on the client, there is only one IP address and default gateway for NIC B-1 and B-2. If they are deleted, neither of NIC B-1 nor B-2 can be used. The reason why we use the Linux bonding driver will be explained in the next subsection.

2) **The reply data from the server must be split into two parts.** Both of the uplinks will transmit packets through WiFi channel. Originally the reply packets from the server would have also been received via WiFi channel. However, the asymmetric VLC link uses VLC channel to receive downlink data. Thus, we must extract some of the reply packets and redirect them to the VLC channel.

One solution for this problem is to separate the reply packets by their different destination IP addresses and redirect the selected ones using the static routing table on the AP. However, because the Linux bonding driver is run under mode 6 on the client, there is only one IP address on the client. This means all of the reply packets sent to the client have the identical destination IP address.

3) **The Linux bonding driver cannot control a wireless NIC.** In our hybrid system, NIC B-1 is a wireless NIC. The Linux bonding driver can only control and bond NICs that support the *ethtool* command. Since wireless NICs do not support *ethtool*, NIC B-1 must be a wired NIC. In order to investigate the performance of the hybrid system under WiFi contention, there must be a wireless link between the AP and client.

One way to keep the wireless link is to add a second AP between NIC B-1 and the original AP. The second AP connects with NIC B-1 via Ethernet cable and connects with the original AP via WiFi channel. Unfortunately, this method does not work because the second AP automatically modifies the Address Resolution Protocol (ARP) packets it receives and ruins the packets redirection. Another approach is to add a wireless NIC, called NIC W-1, on the client. NIC B-1 and NIC B-2 transmit packets to the AP via NIC W-1. The AP sends downlink data to NIC B1 through NIC W-1. One issue here is that NIC B-1 will not deal with the data that are received from NIC W-1. An extra virtual machine under bridge networking mode is built on the client to collect the received data from NIC B-1 and send them back to NIC B-1. However, the virtual bridge vmnet0 that connects the virtual machine and the client is uncontrollable. Thus, the virtual machine cannot redirect the received packets to NIC B-1.

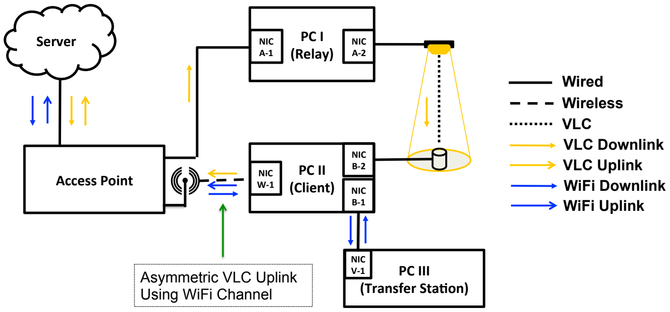

Fig. 4. hybrid system architecture

D. *System Design*

Fig. 4 illustrates the hybrid system architecture of aggregated VLC and WiFi for indoor Internet access.

There are two duplex links in the system. They are combined and controlled by the Linux bonding driver run on the client. The Linux bonding driver schedules the traffic between the two links and thus achieves aggregation of downlink channels. For the asymmetric VLC link, request packets generated by NIC B-2 are captured and transmitted to NIC W-1. Then they are sent through the WiFi uplink and AP, and arrive at the server. Reply packets from the server are forwarded through the AP and relay (PC I), and finally arrive at NIC B-2. For the duplex WiFi link, request packets generated by NIC B-1 are also captured, transmitted to NIC W-1, sent through the WiFi uplink and AP, and arrive at the server. Reply packets from the server are forwarded through the AP and arrive at NIC W-1. The packets are captured from NIC W-1 and sent to NIC B-1. NIC B-1 sends them to NIC V-1 on the transfer station without any manipulation. On the transfer station the packets are captured, adjusted, and put back to NIC V-1. Finally, NIC B-1 receives the packets from NIC V-1. The packet capture, transmission between NICs on the same device, and manipulation are realized by socket programming. The solutions to resolve the issues mentioned in the previous subsection are as follows:

1) To solve the problem mentioned in challenge (1), the Linux bonding driver under mode 6 is used. When the Linux bonding driver is run under mode 6, it will create a logical interface above the NICs on the client. The IP address of the logical interface is the IP address of the client. The MAC address of the logical interface is the MAC address of

the current active NIC, which is working at this time. The logical interface has full control of the NICs on the client and it can switch the active NIC rapidly to achieve bandwidth aggregation. Higher protocol stack layers on the client and other devices (server, AP, and relay) are not aware of the NICs on the client. They communicate with the logical interface on the client. Thus, packets can be received through either NIC B-1 or B-2 because finally they are all received through the logical interface. To avoid the scenario that the client uses NIC W-1 to transmit data, the metric of NIC W-1 is set to be larger than the logic port.

2) To solve the challenge mentioned in challenge (2), we use the MAC addresses of the NICs to split and redirect packets, and make sure the reply packets are received by the corresponding NICs. The Linux bonding driver will send an ARP reply with the MAC address of the chosen NIC to update the next-hop device's ARP table before a TCP connection is built. If the server and the client are in the same subnet, the server's ARP table is updated. Otherwise, only the ARP table of the AP is updated. By modifying the sender hardware address (SHA) of the ARP reply, the server's ARP table is updated to the NIC to which the reply packets will be redirected. If NIC B-1 is chosen to send the first SYN segment, the SHA in the ARP reply is changed to the MAC address of NIC W-1 (Fig. 5(b)). If NIC B-2 is chosen, the SHA in the ARP reply is changed to the MAC address of NIC A-1 (Fig. 6(b)). Thus, the next-hop device sends reply packets with corresponding destination MAC address according to its ARP table. Once the packets reach NIC A-1, the destination MAC address is reverted to the MAC address of NIC B-2 and transmitted to NIC B-2 (Fig. 6(a)). Before this can be realized, the AP must know which port is connected to NIC W-1 and which one is connected to NIC A-1. The AP can update its forwarding table by learning from the packets that transit through each port. Thus, the source MAC addresses of each of the packets sent from NIC B-1 and B-2 are changed to NIC W-1's MAC address (Fig. 5(b), Fig. 6(b)). Notice that if the sender hardware address and the source MAC address of an ARP reply packet is not identical, the sender hardware address will be used to update the next-hop device's ARP table.

3) To solve the issue noted in challenge (3), an extra transfer station is needed. NIC V-1 on the transfer station connects with NIC B-1 via Ethernet cable. NIC B-1 receives data from NIC W-1 and sends them to NIC V-1. NIC V-1 adjusts the destination MAC address of the data and sends them back to NIC B-1. Since the data come from another computer, NIC B-1 will accept and deal with them.

IV. EXPERIMENTS

In the first two experiments, we set the VLC communication distance to 1.5m. Thus, we evaluate the data rate capacity of the VLC downlink to the calculated theoretical upper bound under 1.5m by using Wondershaper [10]. By adding contending WiFi users, we investigate the variation of throughputs experienced by the hybrid system as compared to two conventional systems. These are the asymmetric system and the WiFi-only system. All throughput results are averaged over 10 runs. Each run has a duration of 60 seconds.

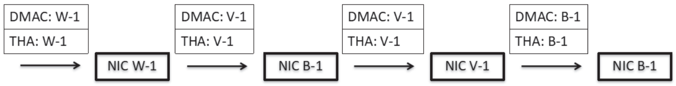

(a) WiFi Downlink

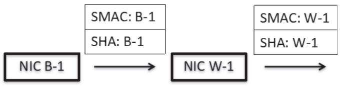

(b) WiFi uplink

Fig. 5. MAC Headers of WiFi Link Flow. a) Downlink. b) Uplink. SMAC is Source MAC address. DMAC is Destination MAC Address. SHA is Sender Hardware Address. THA is Target Hardware Address.

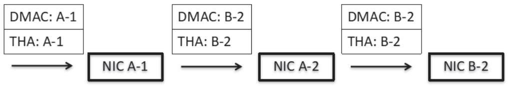

(a) VLC downlink

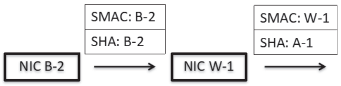

(b) VLC uplink

Fig. 6. MAC Headers of VLC Link Flow.

The RTTs of the WiFi and VLC links of the hybrid system are tested in the third and fourth experiments under the same configuration as the first two experiments. It is necessary to investigate the differences of the RTTs between all the links of a hybrid system when the load-balancing algorithm is designed for the hybrid system. The throughputs and RTTs are measured by *iperf* [11] and *ping* command. The contending users also use *iperf* to generate uplink and downlink traffic. All RTT results are averaged over 10 runs. Each run transmits 10 Internet Control Message Protocol (ICMP) packets.

In the fifth experiment, we evaluate the data rate capacity of the VLC downlink of the hybrid system and the asymmetric system against calculated theoretical upper bounds under different communication distances. We study the variation of the throughput of the hybrid and asymmetric systems with different VLC communication distances.

A. Test Setup

The test setup for the experiments is shown as Fig. 7. The AP is NETGEAR Wireless Dual band Gigabit Router WNDR4500 operating under IEEE 802.11 b/g/n. Both the contending server and test devices are laptops that connect to the AP via wired Ethernet (IEEE 802.3). Two smartphones (iPhone 5s and Google Nexus 5X) and two iPads (iPad mini 2 and iPad 5) are used as contending WiFi devices. The tested systems are the hybrid system, the asymmetric system, and the WiFi-only system. We test one of the three systems each time.

The signal-to-noise ratio (SNR) values of VLC signal under various communication distances (the distances between the

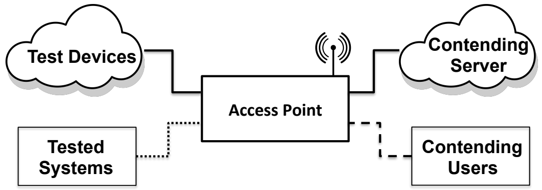

(a) Fixed VLC Communication Distances

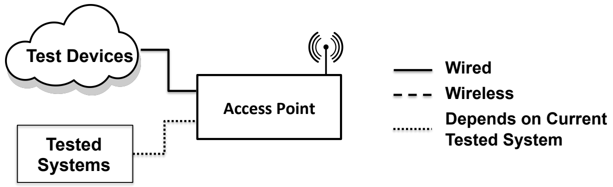

(b) Variable VLC Communication Distances

Fig. 7. Testbed

VLC transmitter and the receiver) are collected in LESA VLC Testbed at Boston University. The transmitter is Cree luminaire (CR22-32L-35K-S) of size $46 * 24cm^2$. The receiver is Thor Labs APD120A2. By using Shannon-Hartley Theorem, the theoretical upper bounds of the VLC channel capacity under different VLC communication distances are calculated. Since TCP is the dominant transport protocol on the Internet, we test the systems with TCP. However, TCP requires that the packets should be received in order. Thus, in a system with multiple links, the system performance can be degraded if there are significantly differences between different links. Here we set the VLC bandwidth to be 10MHz to establish comparable VLC data rates at initial conditions along with a reference WiFi data rate.

B. Results and Analysis

1) Throughput vs. Number of WiFi Contending Users: In the first experiment, we study how the downlink throughput of three systems changes with the number of contending WiFi downlink users. The contending WiFi users, the contending server, and the testbed client are each connected to the same AP. Two laptop test devices are also connected to the same AP with the tested client. As shown in Fig. 8(a), the throughputs of the three systems decrease as the number of contending WiFi downlink users increases. The downlink throughput of the hybrid system is almost the summation of the downlink throughput of the WiFi-only and asymmetric systems. The WiFi-only system decreases faster than the asymmetric system. This result is due to the fact that the contending WiFi downlink traffic affects only the uplink of the asymmetric system, but it affects both the downlink and uplink of the WiFi-only system.

In the second experiment, we investigate how the throughput of the three systems changes with the number of contending WiFi uplink users. Unlike in the first experiment that only has one node (AP) sending TCP packets, there are several uplink users sending a lot of TCP packets in the second experiment. Regardless that the client downloads data through the WiFi channel or the VLC channel, the acknowledgements are transmitted via WiFi channel. When the WiFi channel is congested, the performance of all three systems will be affected. The more users we have, the less bandwidth each single user can get. The throughput of the systems will be degraded when there is insufficient bandwidth for uploading acknowledgements. As Fig. 8(b) shows, the throughputs of all of the three systems decline drastically with the increasing number of the contending uplink users. It is worth noting that the throughput of the hybrid system is even lower than the throughput of the WiFi-only system when there are four contending users. There are two downlink channels in the hybrid system, which leads to more uplink traffic than the WiFi-only system. The more uplink traffic we have, the more total throughput of the network there will be. The total throughput of the network will increase with the total uplink traffic of the users until it reaches the maximum. After that, the total throughput of the network starts decreasing because of

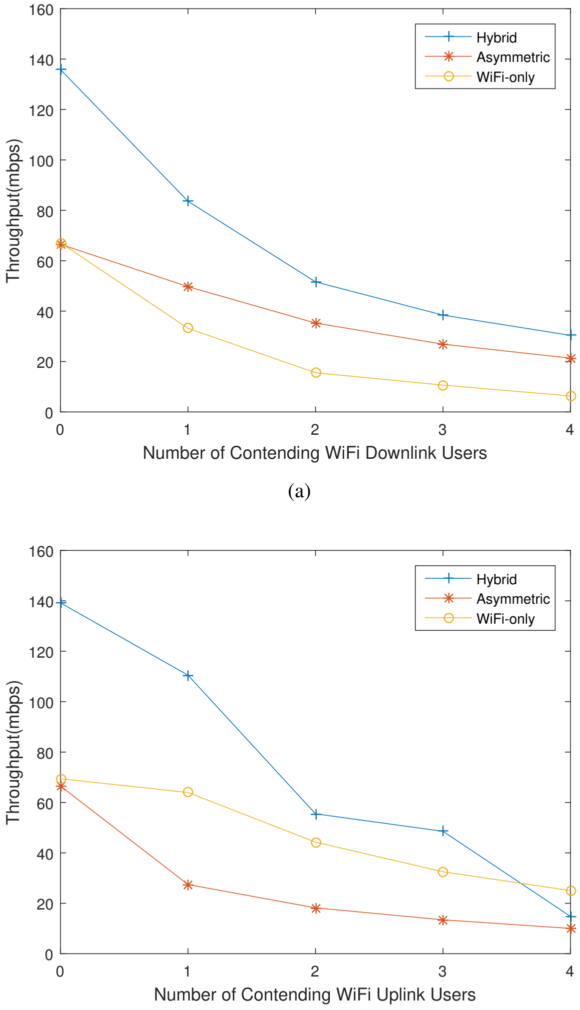

(a)

(b)

Fig. 8. (a) Throughput vs. Number of WiFi Contending Downlink Users. (b) Throughput vs. Number of WiFi Contending Uplink Users.

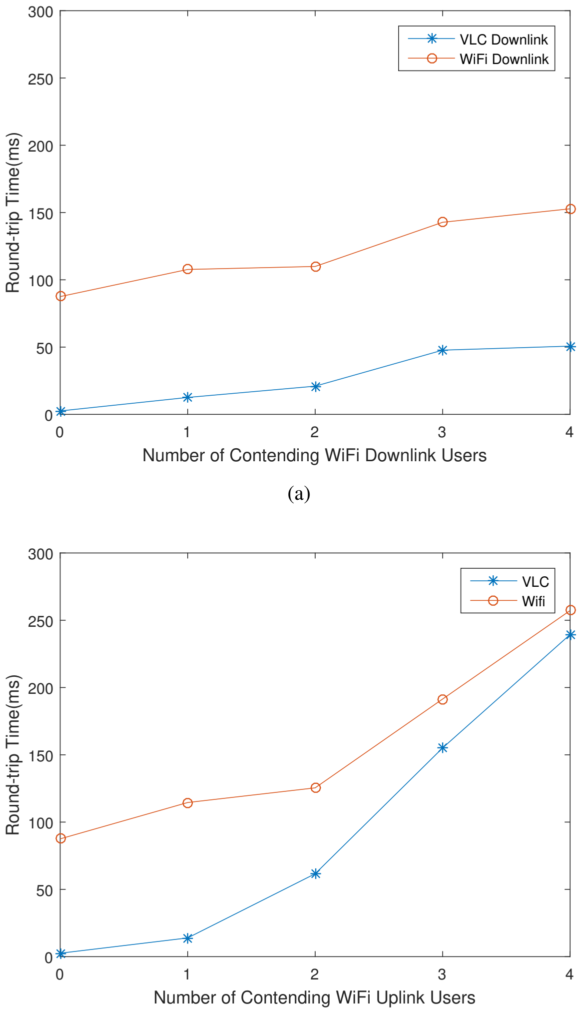

(a)

(b)

Fig. 9. (a) RTT vs. Number of WiFi Contending Downlink Users. (b) RTT vs. Number of WiFi Contending Uplink Users.

collisions [12]. The more users we have, the more remarkable this phenomenon will be. The decrease of the total throughput results in the decrease of the throughput of each user. When there are four contending users and a client of the hybrid system in the network, the total throughput will decline from the maximum and the throughput of a single user will degrade abruptly.

In most of the cases, our hybrid system can achieve aggregated throughput much larger than the asymmetric system or the WiFi-only system. If the user needs a high data rate, the hybrid system can be utilized. However, if the user is in the environment where there are many uplink users, the thresholds of the number of uplink WiFi users and the total load should be considered.

2) RTT vs. Number of WiFi Contending Users: The third and fourth experiments investigate the RTT of the VLC and WiFi links of the hybrid system client. The results demonstrate the need for future work towards load balancing. As Fig. 9(a) and Fig. 9(b) show, the VLC link has less RTT than the WiFi link in each of the scenarios. The obvious difference between the RTTs of two links results from the collisions on the WiFi downlink channel. In a system with multiple interfaces such as our hybrid system, traffic can go through different links with different characteristics, including with different RTTs. Since packets can experience different transmission time, the order of the received packets can be different from their original order. For protocols like TCP, the receiver is responsible for reordering the received packets before use. This reordering creates delays due to waiting on slower links for some of the packets [8], with a corresponding degradation of the average RTT of the system. In order to avoid performance degradation introduced by the large RTT of the slow link in the hybrid system, the variation in RTT among different links should be considered when a load-balancing algorithm is designed. Reference [13] presents an adaptive scheduling algorithm based on different RTT paths. The main idea is to calculate how many packets can be transmitted through the fast path before a packet on the slow path is received. This method can reduce the system latency, which is a function of RTT, but then the fast link is not fully utilized.

As shown in Fig. 9(a), the RTTs of the WiFi link and the VLC link increase at the same speed. The ping command requires little bandwidth to transmit ICMP packets. Even when there are contending downlink users, the ICMP packets can be sent without blocking. The RTTs of both links are only affected by the increase in WiFi uplink traffic. In this scenario, the RTT of the WiFi link is always much larger than the RTT of the VLC link. Latency-sensitive users can utilize the algorithm in [13] to minimize the system latency. Users who want to maximize throughput also need to consider the RTT. Large RTT results in large retransmission timeout (RTO), which leads to long channel idle time and degrades throughput. Thus, full utilization of the available links may not always provide maximum throughput.

If the WiFi uplink channel is congested with contending users, both of the WiFi and VLC links will get large RTTs. As shown in Fig. 9(b), the RTT of VLC link increases faster than the RTT of WiFi link. The reason is that the speed of the WiFi uplink channel is much less than the speed of the WiFi downlink channel. The dominant factor is the contending uplink traffic, which affects both links. The single user will get less uplink performance when there are more contending users. If there is not enough uplink bandwidth for the two links, the downlink channel of the two links will experience significant idle time. In this situation, the RTTs of the WiFi and VLC links will be similar under low uplink rates. In the environment with contending uplink users, the difference of the RTTs between WiFi and VLC links may change significantly with the number of contending users and uplink bandwidth. For latency-sensitive users using the adaptive scheduling algorithm in [13], there is no much difference from the environment with WiFi contending downlink users because the algorithm requires monitoring RTTs of available links. For the users who

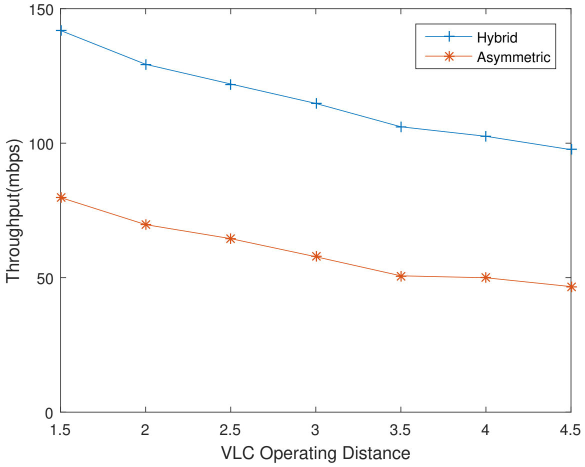

Fig. 10. Throughput vs. VLC Operating Distance

want high throughput, the difference of the RTTs between available links should be considered in real time. With this information, the load-balancing algorithm on the hybrid system client can utilize the available bandwidth efficiently with appropriate RTO.

3) Throughput vs. VLC Operating Distance: Due to the in-order received packets requirement of TCP, the systems with multiple links can experience degraded throughput which is much lower than the throughputs of single VLC link and WiFi link when there are significant differences between the data rates of different links. The data rate of VLC is more sensitive than the data rate of WiFi, which can result in throughput degradation in the proposed hybrid system. In the last experiment, we use Wondershaper to set the data rate of the VLC link to the upper bounds under different VLC operating distances. The throughputs of the hybrid system and the asymmetric system are measured under these upper bounds. Fig. 10 shows that the throughputs of the two systems decline with the same speed until 4.5m. The throughput of the hybrid system is almost the summation of the throughputs of its VLC link and WiFi link, which means that the degradation does not happen. Since VLC has high frequency reuse efficiency, the client of the hybrid system will have access to another VLC access point before the performance declines drastically. This scheme is viable in the dense scenarios that exist indoors.

V. CONCLUSION

In this paper, we design and implement a practical hybrid WiFi-VLC system comprised of two links. One is a duplex WiFi link; the other is an asymmetric duplex WiFi-VLC link. The asymmetric link consists of a VLC downlink and a WiFi uplink. These two duplex links work simultaneously to provide aggregated bandwidth. The system is implemented, and is used to evaluate throughput and RTT for this class of system. Results show that the proposed hybrid system outperforms both a WiFi-only system and an asymmetric system under most scenarios in terms of throughput. With respect to RTT, results for the two links in the hybrid system under congested WiFi show the tradeoff between the bandwidth utilization and latency in algorithm design for the hybrid system. Lastly, varying the VLC link distance in the system demonstrates satisfactory throughput and shows the feasibility of the hybrid system for use in indoor environments.

In the future we anticipate the design of load balancing algorithms that trade off energy efficiency, resource allocation, throughput, and RTT under different operating scenarios. We will explore cross-layer design issues to realize this optimization, with a goal to continously improve performance for the next generation of heterogeneous optical and RF wireless networks.

ACKNOWLEDGMENT

This material is based upon work supported by the National Science Foundation under Grant No. CNS-1617924.

REFERENCES

[1] Zabih Ghassemlooy, Stanislav Zvanovec, Mohammad-Ali Khalighi, Wasiu O. Popoola, and Joaquin Perez. Optical wireless communication systems. *Optik - International Journal for Light and Electron Optics*, 151:1–6, 2017.

[2] Toshihiko Komine and Masao Nakagawa. Fundamental analysis for visible-light communication system using led lights. *IEEE transactions on Consumer Electronics*, 50(1):100–107, 2004.

[3] Sihua Shao and Abdallah Khreishah. Delay analysis of unsaturated heterogeneous omnidirectional-directional small cell wireless networks: The case of rf-vlc coexistence. *IEEE Transactions on Wireless Communications*, 15:8406–8421, 2016.

[4] Sihua Shao, Abdallah Khreishah, Moussa Ayyash, Michael B Rahaim, Hany Elgala, Volker Jungnickel, Dominic Schulz, Thomas DC Little, Jonas Hilt, and Ronald Freund. Design and analysis of a visible-light-communication enhanced wifi system. *Journal of Optical Communications and Networking*, 7:960–973, 2015.

[5] Tarik Borogovac, Michael B Rahaim, Malika Tuganbayeva, and Thomas DC Little. lights-off visible light communications. In *GLOBECOM Workshops, IEEE*, pages 797–801, 2011.

[6] Michael B Rahaim, Anna Maria Vegni, and Thomas DC Little. A hybrid radio frequency and broadcast visible light communication system. In *GLOBECOM Workshops, IEEE*, pages 792–796, 2011.

[7] Thomas Davis, Willy Tarreau, Constantine Gavrilov, Chad N Tindel, Janice Girouard, and Jay Vosburgh. Linux ethernet bonding driver howto. *available as download (bonding. txt) from the Linux Channel Bonding project Web site http://sourceforge. net/projects/bonding/(November 2007)*, 2011.

[8] Karim Habak, Khaled A Harras, and Moustafa Youssef. Bandwidth aggregation techniques in heterogeneous multi-homed devices: A survey. *Computer Networks*, 92:168–188, 2015.

[9] Sihua Shao, Abdallah Khreishah, Michael B Rahaim, Hany Elgala, Moussa Ayyash, Thomas DC Little, and Jie Wu. An indoor hybrid wifi-vlc internet access system. In *MASS, IEEE*, pages 569–574, 2014.

[10] Bert Hubert. The wonder shaper, 2002.

[11] Ajay Tirumala, Feng Qin, Jon Dugan, Jim Ferguson, and Kevin Gibbs. Iperf. 2006.

[12] David Malone, Ken Duffy, and Doug Leith. Modeling the 802.11 distributed coordination function in nonsaturated heterogeneous conditions. *IEEE/ACM Transactions on networking*, 15(1):159–172, 2007.

[13] Farhan H Mirani, Nadia Boukhatem, and Minh Anh Tran. A data-scheduling mechanism for multi-homed mobile terminals with disparate link latencies. In *VTC 2010-Fall, IEEE*, pages 1–5, 2010.