Improving Connectivity, Coverage, and Capacity in 60 GHz Indoor WLANs Using Relays

Swetank Kumar Saha
University at Buffalo, SUNY
swetankk@buffalo.edu

Li Sun
University at Buffalo, SUNY
lsun3@buffalo.edu

Dimitrios Koutsonikolas
University at Buffalo, SUNY
dimitrio@buffalo.edu

ABSTRACT

60 GHz technology, standardized by IEEE 802.11ad, has emerged as an alternative to 2.4/5 GHz legacy WiFi for building multi-Gigabit general purpose home/enterprise WLANs. However, limited range and high vulnerability to blockage of millimeter-wave (mmWave) links can severely affect WLAN performance. These problems coupled with long re-connection time for a broken link pose a major challenge to the realization of 60 GHz WLANs. In this work, we look at the *relay* architecture defined in the 802.11ad standard as a solution to these issues and propose extensions to it for achieving gains in both WLAN coverage and capacity.

Categories and Subject Descriptors

C.2.1 [**Network Architecture and Design**]: Wireless communication

Keywords

60 GHz; 802.11ad; Relays

1. INTRODUCTION

IEEE 802.11ad-based WLANs can potentially provide multi-Gigabit throughput with a maximum bit-rate of 6.76 Gbps by leveraging ultra-wide channels (2.16 GHz). However, 60 GHz links, owing to the small wavelength of the carrier wave, suffer from high attenuation and cannot easily penetrate solid objects. As a result, 60 GHz links are short and cannot sustain human blockage without severe link quality degradation. To overcome these challenges, we turn our attention to the use of relays to augment the existing Access Points (APs) in the WLAN.

First, we experimentally demonstrate the high vulnerability of 60 GHz links to human blockage and show that relays can naturally help alleviate the *connectivity* problem. Then, we discuss the possibility of extending the relay functionality as described in the 802.11ad standard[10] and propose a new WLAN architecture to improve *coverage* and increase network *capacity* with concurrent transmissions. We discuss the research challenges involved in each of the above scenarios, which we plan to address in our future work.

S3'15, September 11, 2015, Paris, France.

ACM ISBN 978-1-4503-3701-4/15/09 ...$15.00.
DOI: http://dx.doi.org/10.1145/2801694.2804446.

Related Work A number of works have considered the use of relays in 60 GHz wireless networks and studied relay selection [8, 2], MAC design [6], interference-aware link scheduling [5, 4], and relay placement [9]. All these works consider WPANs, where all clients are located in a single room, the traffic pattern consists mainly of client-to-client connections, clients themselves can act as relays, and each client maintains a connection to the AP. In contrast, we consider a much more challenging scenario; an enterprise WLAN consisting of multiple offices, corridors, etc, with multiple APs and dedicated relays deployed to provide coverage to a large number of clients. [8, 9] considers the use of dedicated relays but assume a uniform grid placement which in not practical. Further, most of these works [8, 5, 4, 9] study the problem only from a theoretical perspective, without considering protocol design. Other works design protocols that make unrealistic assumptions, e.g., use of omni-directional antennas [2] or quasi-omni beam patterns [6], or a uniform/grid relay placement [9]. Finally, some works have shown the benefits of relays through simulations [3, 1]. To our best knowledge, this the first work that demonstrates the potential benefits of relays with COTS devices.

2. EXPERIMENTS

To quantify the effect of human blockage/mobility, we conducted experiments with COTS 802.11ad hardware in a Hall inside the Computer Science building of the University at Buffalo.

2.1 Experimental Setup

We setup 802.11ad link using a Dell Latitude E420 laptop equipped with a Wilocity wil6210 802.11ad radio and a Dell Wireless Dock D5000. The dock has an 802.11ad wireless interface and acts as an AP. Another laptop is connected to the dock through a Gigabit Ethernet interface to generate/receive TCP traffic. The use of the Ethernet interface limits the throughput in our experiments to 1 Gbps, even though the wireless link itself is capable of much higher speeds. The Wilocity radios do not allow us to control the PHY layer data rate and use their own rate adaptation algorithm[1] and an in-built beamforming mechanism. They only report the current PHY data rate and an RSSI value between 0 and 100. We use the average PHY data rate as an indicator of the maximum achievable throughput.

2.2 Experimental Methodology

All the reported results are the average of 5 10-second iperf sessions. We experiment with different link quality (as indicated by RSSI) by varying the distance between the Transmitter (Tx) and Receiver (Rx). The Tx is placed at a height of 5'6" and Rx at 3'. We

[1]The following PHY data rates are supported by our cards (in Mbps): 385, 770, 1155, 1540, 1925, 2310, 3080, 3850.

found that a higher Tx placement, in general, gives better performance.

For each distance, we considered two kinds of blockage: static and mobile. The static case represents a scenario where the LOS link is disrupted permanently. For the static case, we repeat the experiment with two sub-cases. In the first one (Static 1), human blockage happens mid-way between the Tx and Rx. In the second case (Static 2), the human stands close to the Rx. We did not consider the case where the blockage is close to the Tx, as in practice APs will be placed much higher than the average human height. In the mobile case, a person walks in a random fashion along the LOS path between the Tx and Rx. Note that in this case the LOS link is disrupted only intermittently.

We place the relay at the same height as the Tx. The relay is placed at 1/4 (Pos 1), 1/2 (Pos 2), or 3/4 (Pos 3) of the distance between the Tx and Rx but laterally displaced by 3 ft from the straight line joining the Tx and Rx, so as to form a triangular topology. The relay path consists of two links: one from the Tx to the relay ($link_1$) and other from the relay to the Rx ($link_2$). We do measurements for each of these links individually and then calculate the average PHY data rate for the relay path as $1/[(1/r_1) + (1/r_2)]$, where r_1 and r_2 are the average PHY data rates of $link_1$ and $link_2$ respectively. Note that the reported results assume the relay is half-duplex. The 802.11ad standard allows for full-duplex relays as well, in which case the performance gain will be higher.

2.3 Results

Figure 1 plots the average PHY data rate, for 7 cases: *Base* (no blockage), *Static 1*, *Static 2* and *Mobile* blockage, as described in the previous section, and *Relay* (for each of three relay positions). For small separation between the Tx and Rx (8 ft), all three cases of blockage cause complete disconnection. In fact, the Tx even fails to find an alternative NLOS path to the Rx and the link cannot recover unless the blockage is removed. This presents the best application case for relays where without relay, communication would not be feasible at all. Also notice that the relay's average PHY data rate is much greater than half of the base case. This is because each of the two links involved in the relay are faster than the direct link. For higher separation of 16, and 24 ft, blockage causes link degradation but does not cause disconnection. For these intermediate distances, at least one of the relay positions offer a gain (although marginal) over the blockage scenarios. In such cases, if the switching delay was taken into account the total gain of using a relay may in fact be negative. This highlights an important fact that online channel measurements/estimation may be required before switching to the relayed path. When we further increase the separation to 32 ft, link degradation becomes severe and using the relay has a clear advantage over the degraded link, with the relayed link offering upto 1250 Mbps (for Pos 1) higher data rate compared to each of the blockage scenarios. Even in the worst case, the relayed link provides a gain of 650 Mbps (for Pos 2). Lastly, for larger separations of 40 and 48 ft, blockage causes complete disconnection and the use of relay allows for PHY data rates of upto 1400 Mbps. Overall, relays are most useful for very short and long distances but provide moderate to no gain for intermediate distances.

Interestingly, all three relay positions perform rather similarly, with their average PHY data rates for a given distance being almost equal (except for 32 ft). Note that the average PHY data rate of the relay path is affected by both $link_1$ and $link_2$. Nonetheless, this indicates that relay performance is not highly dependent on its placement and in case of blockage a client can achieve performance gains by connecting to any of the visible relays.

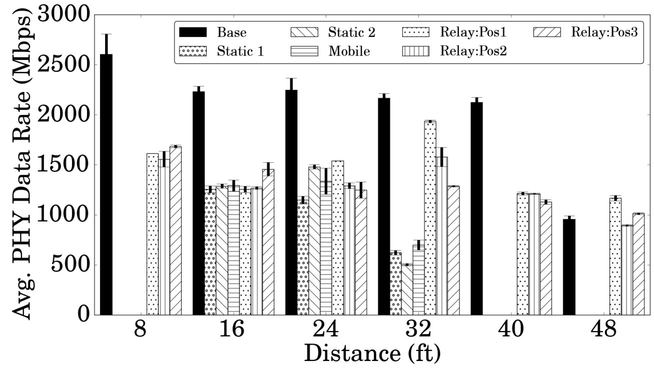

Figure 1: Avg. PHY data rate as a function of distance

3. RELAY-BASED WLAN ARCHITECTURE

In the previous section, we showed cases where relays can help improve performance when the LOS link experiences permanent or temporary blockage. In this section, we propose a relay-based 60 GHz WLAN architecture. We discuss how we can employ the relay functionality as described in the 802.11ad standard document to overcome blockage and through modifications, use it as a mechanism to improve overall WLAN coverage and network capacity. We first give a summary of the relay functionality as envisioned in the 802.11ad standard and then describe our proposed architecture.

3.1 Relays in 802.11ad

The 802.11ad standard proposes an architecture where wireless terminals (clients) can be used as relays between two communicating stations (APs or clients). Stations during association indicate if they are willing to act as relays. The standard requires a one-time relay identification procedure for a pair of communicating stations to discover possible relays. Clients request the AP for a list of relay nodes in the network. The AP responds with a list of possible relay stations and reserves two service periods (SPs) for each relay in the list. One SP is used for beam-forming training between the source and the relay and the other one between the relay and the destination. Once the beamforming weights are learned by all stations involved, a link/path measurement is done through each of relays identified and of the direct path. The standard leaves it to implementation to decide how to choose the best relay. The standard allows for multiple modes of operation, where a source-destination pair may decide to use a relay or communicate directly or use both paths simultaneously, if full-duplex operation is supported. The standard also defines fast switching procedures between the relayed and the direct path.

Note that 802.11ad assumes the co-operation and availability of stations/clients in the network to act as relays. In a WLAN, clients may not want to participate in the relay setup due to energy concerns. Further, the standard allows for link quality measurements over different relays only once during the relay initialization.

3.2 Proposed Architecture

We propose an enterprise WLAN architecture where APs are augmented with secondary APs acting as *relays*. The primary APs, in addition to the 802.11ad radio, have a wired connection to the rest of the backbone network, whereas secondary APs/relays only host an 802.11ad wireless interface. The exact functions of the secondary APs/relays vary based on the use-case (connectivity improvement, coverage increase, or capacity addition). We discuss each of these cases separately in the sub-sections below.

Note that, in contrast to a WPAN architecture considered in most previous works, where the majority of the traffic flows is between clients, in our scenario, communication always takes places between a client and the wired network via one of the primary APs, as in legacy WLANs.

We believe that such an architecture allows for incremental deployment, where secondary APs can be simply added to the WLAN in an ad-hoc fashion and discovered by the primary APs. Given the fact that an 802.11ad WLAN will always require a much denser deployment of primary APs, compared to its 802.11n/ac counterpart, having such an architecture can save on both manual labor and cost of deployment by reducing the number of primary APs required and the amount of wiring involved in connecting them to the rest of the network.

3.3 Relay Placement

Consider a scenario where an enterprise network upgrade replaces all existing 802.11n/ac APs with 802.11ad APs. These APs will be augmented by a number of relays to (i) cover "coverage holes" created due to the limited range of 60 GHz links compared to 2.4/5 GHz links and (ii) to provide back up paths in case of link outages due to human blockage. We model the placement of relays as a variation of the *Museum Problem*, where the APs/relays play the role of the guards and the goal is to ensure that X% of a given floorplan area can be "seen" by at least m guards. We consider two cases: (i) for a given number of relays, find the optimal placement to maximize X and (ii) for a target X, find the minimum number of required relays and their placement. Note that for a theoretical treatment of the problem, APs and relays are considered identical. In practice, the type of functionality implemented on the relays (e.g., whether they transmit beacons or not) can bring different benefits in terms of connectivity, coverage, and capacity improvements, each of which we discuss next.

3.4 Improving Connectivity

The 802.11ad relay-based architecture described above can naturally act as a mechanism to combat the blockage problem and improve connectivity. If a source station experiences blockage in the direct path to the destination, it selects the relayed path. Note that this functionality does not require any changes in the standard. The secondary APs behave just like clients, except that they do not generate any traffic of their own and they are not the final destination of any traffic generated by primary APs or other stations in the WLAN.

In the context of improving connectivity, this architecture brings up interesting research questions.
(1) How can we enable frequent measurement of path/link quality through different relays between a source and destination, without excessive overhead?
(2) The standard assigns *one* relay to each link. How can we allow for multiple secondary APs to act as possible relay candidates for a given link so that the link can switch relays online based on measured link qualities?

3.5 Increasing Coverage

The severe path loss of the 60 GHz carrier-wave with increasing distance limits the coverage of an 802.11ad AP compared to its 802.11n/ac counterpart. For a WLAN to provide a good user-experience, the AP discovery and re-connection time in case of device mobility need to be kept low. Recently, [7] reported that AP discovery latency ranges from 5 ms to 1.8 s for static clients and upto 12.9 s for mobile clients.

To overcome this problem, we propose modifications to the relay architecture which allow to increase the WLAN coverage without opting for a highly dense AP deployment. Allowing secondary APs/relays to act as proxy APs by enabling beacon transmissions and processing of association requests from clients automatically extends the coverage of each individual primary AP and in turn of the whole WLAN. Note that the relays are themselves associated to an AP. The AP can now allot additional time (e.g., an extra SP in its TDMA or polling-based MAC [10]) to the relay node to allow for communication with the clients associated to the relay terminal.

In the context of improving coverage, we consider the following research directions.
(1) Each relay can now serve more than one clients and needs its own channel allocation mechanism to distribute the time allotted by the AP to those clients. How can we seamlessly integrate the schedule of an AP and the schedules of each the secondary APs/relays connected to that AP?
(2) How do we allow for interference free beaconing by the primary AP and its secondary APs/relays?

3.6 Adding Capacity

The combination of polling-TDMA MAC and the spatial reuse due to the directional transmissions employed by 802.11ad allows for opportunities to further expand the capacity of 60GHz WLANs by enabling concurrent transmissions in a single SP. The presence of relays opens up the possibility of having even more concurrency by allowing an otherwise idle relay (assigned to another client-AP link) to service another client in parallel with its primary AP. Although the relay still needs another SP to transmit the packets to the AP, faster links between the client and relay, and the relay and AP can provide an overall gain for the network. Here again, there are multiple issues to address.
(1) We first require a framework that allows the AP to build the network conflict graph, compute an optimal TDMA schedule, and distribute it to the relays and clients.
(2) In a given SP, a secondary AP can either be assigned to a primary AP-client link (offering a backup path in case of link blockage) or serve another client concurrently with a scheduled AP-client link. How to make this decision for each relay such that a given objective (e.g. throughput, reliability) is maximized?

4. REFERENCES

[1] Z. Genc, G. Olcer, E. Onur, and I. Niemegeers. Improving 60 ghz indoor connectivity with relaying. In *Communications (ICC), 2010 IEEE International Conference on*, pages 1–6, May 2010.

[2] Z. Lan, C.-S. Sum, J. Wang, T. Baykas, F. Kojima, H. Nakase, and H. Harada. Relay with Deflection Routing for Effective Throughput Improvement in Gbps Millimeter-Wave WPAN Systems. *IEEE Journal on Selected Areas in Communication (JSAC)*, 27(8), 2009.

[3] C. S. Leong, B. S. Lee, A. R. Nix, and P. Strauch. A Robust 60GHz Wireless Network with Parallel Relaying. In *Proc. of IEEE ICC*, 2004.

[4] Y. Niu, Y. Li, D. Jin, L. Su, and D. Wu. Blockage Robust and Efficient Scheduling for Directional mmWave WPANs. *IEEE Transactions on Vehicular Technology*, 64(2), 2015.

[5] J. Qiao, L. X. Cai, X. S. Shen, and J. W. Mark. Enabling Multi-Hop Concurrent Transmissions in 60 GHz Wireless Personal Area Networks. *IEEE Transactions on Wireless Communications*, 10(11), November 2011.

[6] S. Singh, F. Ziliotto, U. Madhow, E. M. Belding, and M. Rodwell. Blockage and directivity in 60 ghz wireless personal area networks: from cross-layer model to multihop mac design. *IEEE Journal on Selected Areas in Communications (JSAC)*, 27(8):1400–1413, October 2009.

[7] S. Sur, V. Venkateswaran, X. Zhang, and P. Ramanathan. 60 GHz Indoor Networking through Flexible Beams: A Link-Level Profiling. In *Proc. of ACM SIGMETRICS*, 2015.

[8] C. Yiu and S. Singh. Link Selection for Point-to-Point 60GHz Networks. In *Proc. of IEEE ICC*, 2010.

[9] G. Zheng, C. Hua, R. Zheng, and Q. Wang. Toward Robust Relay Placement in 60GHz mmWave Wireless Personal Area Networks with Directional Antenna. *Accepted in IEEE Transactions on Mobile Computing*, 2015.

[10] IEEE 802.11 Task Group AD. `http://www.ieee802.org/11/Reports/tgad_update.htm`.