Understanding IEEE 802.11n Multi-hop Communication in Wireless Networks

Jan Friedrich, Simon Frohn, Sascha Gübner, and Christoph Lindemann
Department of Computer Science
University of Leipzig
Johannisgasse 26, 04103 Leipzig, Germany

{friedrich, frohn, guebner, cl}@rvs.informatik.uni-leipzig.de

Abstract — **In this paper, we present a measurement study of the multi-hop behavior of the new IEEE 802.11n standard in an indoor mesh testbed. As main contribution, we quantitatively describe characteristics of IEEE 802.11n on multi-hop paths like throughput, aggregate size and utilized MIMO features. Furthermore, we show how to build an IEEE 802.11n indoor testbed capable of multi-hop communication. In an 8-hop chain topology, we observe that channel bonding nearly doubles the throughput for any fixed path length. The mean aggregate size in number of frames at each node is also doubled by channel bonding. The aggregate size decreases with increasing path length. Limiting the aggregate size severely impacts throughput both on single and multi-hop paths. The throughput degrades as the path length is increased as in legacy IEEE 802.11a/b/g. We believe that these findings will enhance the understanding of the performance of IEEE 802.11n in multi-hop networks.**

Keywords: Wireless Mesh Networks; MIMO Measurement Study; IEEE 802.11n Testbeds

I. INTRODUCTION

With the emergence of the IEEE 802.11 technology, wireless mesh networks have experienced a wide-spread deployment drawing substantial interest in academia and industry. Mesh networks can be built with relatively low infrastructure expenditure compared to wired broadband. Therefore, they are particularly attractive for providing fast and cost-efficient coverage for hard-to-wire areas. Particular application areas include disaster scenarios, wireless machine-to-machine communication, and wireless video surveillance.

The IEEE 802.11n standard [7] is the first IEEE 802.11 standard to introduce a MIMO-based physical layer, providing higher data rates up to 600 Mbit/s and higher range and interference tolerance. These features make IEEE 802.11n a promising technology for building carrier grade wireless mesh networks. The main innovations, beside others, responsible for the advantages of IEEE 802.11n are channel bonding, spatial division multiplexing and space-time block coding. Channel bonding doubles the physical data rate by using two adjacent IEEE 802.11a or IEEE 802.11b/g channels, respectively, covering 40 MHz instead of 20 MHz. Spatial division multiplexing and space-time block coding are features that exploit the multi-path propagation of the channel by employing MIMO. In spatial division multiplexing mode the sender transmits different data streams simultaneously along multiple antennas. In space-time block coding mode the sender uses a special coding to transmit a data stream over multiple antennas. The advantage is an increased tolerance of interference and higher transmission range. The high data rates provided by the IEEE 802.11n physical layer can only by harnessed at upper layers if medium access is efficient. Therefore, IEEE 802.11n introduces frame aggregation. Using MPDU (MAC protocol data unit) aggregation multiple frames can be transmitted in an aggregated frame, with the overhead for medium access arising only once.

In this paper, we present a comprehensive measurement study of the multi-hop behavior of the new IEEE 802.11n standard in an indoor mesh testbed. Opposed to previous work [15], [18], and [20], we set up an 802.11n wireless testbed in ad-hoc mode rather than in infrastructure mode, hence, being able to investigate multi-hop communication. We outline the kernel and network system software enhancements to enable multi-hop communication in 802.11n. The presented measurement study using the mesh testbed quantitatively describes characteristics of IEEE 802.11n on multi-hop paths like throughput, aggregate size and utilized MIMO features. We analyze the throughput behavior and its dependence on path length, maximum aggregate size and channel bonding option. We examine the standard Linux rate adaptation algorithm, a crucial element of the IEEE 802.11n efficiency, and investigate its performance in multi-hop scenarios. Furthermore, we reveal details on the aggregate size in a multi-hop flow and its dependence on path length, node position within a flow and channel bonding option. Our main findings are as follows:

- channel bonding nearly doubles the throughput for any fixed path length
- the mean aggregate size in number of frames at each node is also doubled by channel bonding
- mean aggregate size in number of frames at each node decreases with increasing path length
- limiting the aggregate size severely impacts throughput performance for both single-hop communication and multi-hop path
- the advantage of spatial division multiplexing fades away with increasing path length
- throughput degrades as the path length is increased like in legacy IEEE 802.11a/b/g

We believe that these findings will enhance the understanding of the performance of IEEE 802.11n in multi-hop networks. These results are considerably more trustful than corresponding simulation results because they are obtained by measurements in a real-world indoor testbed.

The remainder of this paper is organized as follows. Section II summarizes related work on measurement studies

in IEEE 802.11b/g and IEEE 802.11n networks. In Section III, we introduce our indoor MIMO mesh testbed and outline the technical issues to be solved for enabling multi-hop communication over IEEE 802.11n networks. Section IV presents and discusses measurement results. Finally, concluding remarks are given.

II. RELATED WORK

Recently, LaCurts et al. [12] analyzed traces gathered from 110 different wireless mesh networks deployed by Meraki using both 802.11b/g and 802.11n devices. They studied accuracy of SNR-based bit rate adaptation, the impact of opportunistic routing and the prevalence of hidden terminals. Opposed to this work, we focus on the impact of frame aggregation, spatial division multiplexing and space-time block coding to network performance. Furthermore, we consider an indoor mesh testbed with little interference from 802.11a/b/g background traffic. Kim et al. [10] proposed a modification of the IEEE 802.11 MAC to allow aggregation of unicast and broadcast frames and evaluated it using a wireless node prototype. Opposed to [10], we study frame aggregation in the existing IEEE 802.11n standard using commodity hardware instead of a self developed system.

Halperin et al. [6] showed that wireless packet delivery can be accurately predicted using 802.11n channel state information measurements as input to an OFDM receiver model. Khattab et al. [9] experimentally showed that 802.11n medium access worsens flow starvation as compared to 802.11a/b/g and designed an asynchronous MIMO MAC protocol that tackles the problem. Pefkianakis et al. [15] studied MIMO based rate adaptation in 802.11n wireless networks in a real testbed in infrastructure mode and proposed a MIMO aware rate adaptation scheme. Opposed to [6], [9] and [15], we consider multi-hop communication under 802.11n instead of 1-hop communication in infrastructure mode.

Pelechrinis et al. [16], [17], [18] conducted experimental studies on the behavior of MIMO links in different topologies. They mainly focus on throughput in isolation and with competing 802.11g-links [16], impact of the different 802.11n specific features on the peak performance [17], and packet delivery ratio under different physical data rates [18]. Shrivastava et al. [20] studied the impact of channel bonding and interference of 802.11g on 802.11n-links in a real testbed deployment. Opposed to [16], [17], [18] and [20], we focus on frame aggregation in a multi-hop mesh network instead of a 1-hop infrastructure mode WLAN.

Koivunen et al. [11] presented sample results from a measurement campaign of multi-link MIMO channels at 5.3 GHz in an indoor office environment. Piazza et al. [19] demonstrated a new reconfigurable antenna array for MIMO communication systems that improves link capacity in closely spaced antenna arrays. Opposed to [11] and [19] we focus on MAC mechanisms in IEEE 802.11n rather than on physical layer issues.

In own previous work [5], we characterized the effective throughput for multi-hop paths in IEEE 802.11n wireless mesh networks as a function of bit error rate, aggregation level, and path length. Li et al. [13] proposed an analytical model assuming saturated traffic. They derived the effective throughput and optimal frame and fragment sizes for single-hop links. Papathanasiou et al. [14] investigated through simulations the efficiency of multicast beamforming optimization over IEEE 802.11n WLAN. Opposed to [5], [13] and [14], the findings presented in this paper are derived from measurements in a real 802.11n indoor mesh testbed.

Bicket et al. [3] conducted a comprehensive measurement study in a 37-node 802.11b outdoor mesh network. In [4] we presented ScaleMesh, a miniaturized dual-radio wireless mesh testbed based on IEEE 802.11b/g and conducted a comprehensive measurement study. Opposed to [3] and [4], we focus in this work on an IEEE 802.11n mesh network.

III. INDOOR MIMO MESH TESTBED

A. Background on IEEE 802.11n

The IEEE 802.11n standard features several enhancements for higher throughput in 802.11 wireless communication. Among others, the main features include channel bonding, the utilization of spatial division multiplexing and space-time block coding, and frame aggregation. Channel bonding allows an IEEE 802.11n transmitter to use two adjacent channels in parallel, thus doubling the effective physical data rate. While data rate increases, the communication becomes more sensitive to interference, especially from legacy 802.11a/b/g links, thus collisions and erroneous transmissions occur more frequently. This is why this feature can be enabled manually by the user.

Two MIMO-specific innovations of IEEE 802.11n are spatial division multiplexing and space-time block coding. The first one increases the physical data rate by sending two or more independent data streams in parallel while the last one enhances the robustness of transmission using a special coding, e.g. Alamouti coding [1] for two antennas, to transmit a data stream over multiple antennas and multiple time slots. The choice whether spatial division multiplexing or space-time block coding is employed is encoded in the modulation and coding scheme (MCS) that the transmitter uses. The Atheros chipset, which we employed, is able to use 2x2 MIMO. MCS classes 0 to 7 use one data stream while classes 8 to 15 use two data streams encoded with spatial division multiplexing. According to the notation in [15], we denote the first one single-stream mode and the second one double-stream mode. Furthermore, the MCS class determines the physical modulation and forward error correction scheme, making 0 and 8 the most and 7 and 15 the least robust, respectively.

To fully utilize the high data rates provided by the physical layer, another innovation can be found in IEEE 802.11n. Frame aggregation on the MAC layer allows combining of several data frames into one larger aggregated frame with the overhead for medium access arising only once. Each correctly received frame can be acknowledged by the receiver in a block acknowledgement, reducing the number of frames to be retransmitted. For our hardware the maximum amount of frames to be aggregated is limited to 32 frames. We denote this maximum allowed number of frames to be aggregated as maximum aggregate size.

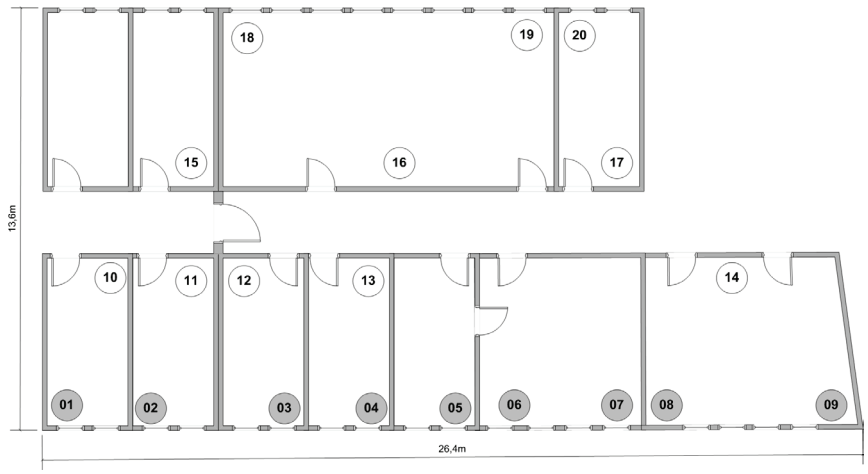

Figure 1. Indoor MIMO Mesh Testbed

Figure 2. 8-hop chain topology

B. Indoor MIMO Mesh Testbed Setup

Our Indoor MIMO Mesh Testbed comprises 20 wireless mesh nodes located in 10 rooms in the department building covering roughly 250 m². The rooms are separated by 15cm thick light-gypsum walls, except for a solid firewall between nodes 02 and 03. An overview of the testbed with the node locations is depicted in Figure 1. Note, that the doors where mainly closed during experiments. Each node consists of a Siemens ESPRIMO P2510 PC with an Intel Celeron 3.2 GHz processor, 512 MB RAM, 80 GB HDD and a D-Link DWA-547 wireless PCI network interface card (NIC). This NIC is equipped with three 5dBi omnidirectional antennas and an AR 9223 Atheros chipset, able to support 802.11n-based MIMO communication in the 2.4 GHz band. Each node runs openSUSE 11.2 as operating system with a modified kernel based on version 2.6.34.

As device driver for the wireless NIC, we employ the *ath9k* driver for Atheros chipsets. To allow remote management of the nodes, each node possesses a Gigabit Ethernet NIC, connected to the subnet of the University of Leipzig through a Gigabit switch. Hence, wireless experiments can be managed from a remote computer and traces can be copied and evaluated through the wired network. Table I shows a detailed description of hardware and software components of the testbed.

TABLE I. TESTBED OVERVIEW

Component	***Description***
PC	Siemens ESPRIMO P2510 Celeron 3.2 GHz, 512 MB RAM, 80 GB HDD
Wireless Card	D-Link DWA-547 PCI NIC equipped with 3 antennas
Chipset	Atheros AR 9223, operating at 2.4 GHz
Operating System	openSUSE 11.2 with kernel version 2.6.34

C. Enabling Multi-hop Communication in IEEE 802.11n

To run our MIMO testbed as a mesh in ad-hoc mode allowing multi-hop communication, we had to modify and extend the existing Linux kernel in several ways.

First to let each wireless node know its neighbors' 802.11n capabilities, the periodically transmitted IBSS beacons had to be extended to carry this information. Issues concerning the joining of nodes to an ad-hoc network had to be resolved, so that finally each node communicates with the offered high data rates. We needed to set up our cards both as normal interface and as a special monitoring interface that allows capturing management and erroneous frames. Furthermore, we changed the *ath9k* device driver to allow a limitation of the number of frames to be aggregated. In addition, we implemented an extensive tracing module that logs the MAC sequence numbers of each frame per aggregate and which ones were erroneous and had to be retransmitted. This data log was later used to map frames on the receiver side with the appropriate frames on the transmitter. We assured that the CPU and IO overhead arising by the tracing do not effect IEEE 802.11n operation by conducting various validation experiments.

Additionally, we improved the performance of the currently implemented aggregation mechanism by allowing single frames to be aggregated when they have to be retransmitted. Note, that this is in conjunction with the IEEE 802.11n standard; however experiments proved that the throughput is increased.

IV. MEASUREMENT RESULTS

A. Experimental setup

For our experimental measurement study, we built up an 8-hop chain topology using the shaded nodes of our indoor MIMO mesh testbed depicted in Figure 1. We positioned the

nodes to let the antennas face into the building to enrich the multi-path scattering, crucial for spatial division multiplexing. We configured all nodes to run in ad-hoc mode and assigned static routing between them. We conducted fixed length measurement experiments starting from an initially idle system. Each experiment lasted 60 seconds. We conducted 10 independent replicates of each measurement experiment and derived the considered performance measures with 95% confidence level. The width of the confidence intervals is depicted as bars in the plots. We used the bandwidth measurement tool *iperf* for Linux [8] to create saturated UDP traffic at the sender with a payload size of 1460 bytes. We chose UDP traffic to limit the influence of the TCP exponential backoff mechanism that may degrade throughput on multi-hop paths. Because our devices run in the 2.4 GHz band we estimated the impact of interference. Therefore, we ran a one-week long-term experiment measuring the throughput on a sensitive link to identify time slots with the least external interference. We noticed that during the working hours between 8am and 8pm, the measured throughput was influenced by external interference, especially due to students who access the web wirelessly through their IEEE 802.11 equipped laptops. Hence, experiments were conducted at night or during the weekend at which time little interference due to 802.11a/b/g background traffic occurred.

Furthermore, we need a bit rate adaptation algorithm to choose the most appropriate MCS class for the current topology and channel conditions. Therefore, we employ the *Minstrel HT* rate adaptation algorithm implemented in the Linux kernel. *Minstrel HT* is the default rate adaptation algorithm for 802.11n in Linux and an advancement of the widely used SampleRate algorithm [2] by Bicket.

B. Measurements in Multi-hop IEEE 802.11n

As an initial experiment, we measure the achievable throughput on a multi-hop chain where we varied the path length from 1 to 8 hops. An example for an 8-hop chain is depicted in Figure 2. We first let the rate adaptation algorithm choose from all supported MCS classes (0 to 15 in our case) allowing to utilize spatial division multiplexing on links with adequate link quality. We repeat the experiment with both activated and deactivated channel bonding option. The measured throughput results are depicted in Figure 4.

To show how the rate adaptation algorithm performs, we plot in Figure 3 the probability mass function of the chosen MCS classes of each transmitter on an 8-hop chain. We see that nearly all rates, both with single-stream and double-stream communication, where utilized. This is a clue that the higher double-stream rates can't be employed on every link. Furthermore, we notice that the rate adaptation algorithm tends more towards single-stream rates if channel bonding is activated. It chooses MCS 5 most frequently while MCS 13 is most commonly used if channel bonding is not activated. We suppose that this is due to the higher interference sensitivity of activated channel bonding, so as a trade-off more robust single-stream rates are chosen. In the next experiment, we restrict the rate adaptation algorithm to only use single-stream rates and quantify the effect on the measured throughput.

We observe in Figure 4 that for 1-hop activated channel bonding raises throughput about nearly 100% with 140 Mbit/s compared to 75 Mbit/s without this option. These results agree with corresponding results of earlier work [15], [17], and [20]. In fact, the quantitative results of Figure 4 lie in between the smallest and largest values of corresponding 1-hop throughput results reported in [15], [17], and [20]. Furthermore, we notice that the throughput degrades with increased path length like in 802.11a/b/g networks and that channel bonding still enlarges throughput significantly for larger path lengths, increasing the throughput on an 8-hop chain about 80% from 4 Mbit/s to 7.3 Mbit/s.

We observe in Figure 5 that limiting the MCS choice severely degrades throughput about 60% for 1-hop flows both for activated and deactivated channel bonding option. We also observe that this effect flattens when path length increases, leading to nearly the same throughput at 8 hops. The reason is on the one hand that more links with lower quality are involved in the multi-hop communication generating a throughput bottleneck. On the other hand, increased medium contention and higher collision probability are limiting factors on longer paths and exceed the effect of higher rate choices. One can expect that this also has an effect for competing flows, as the contention is comparable. Due to space limitations we leave this problem for future work.

To get more insights on frame aggregation under IEEE 802.11n, we plot the mean number of aggregated frames for varying path lengths in Figure 6. We observe that the mean aggregate size with activated channel bonding is nearly twice as much as without this option, about 29 frames compared to 16 frames per aggregate, respectively. We also notice that the aggregate size decreases with higher path lengths and nearly halves on an 8-hop path both for activated and deactivated channel bonding. We suppose that both effects are a result of the increased physical data rate provided with channel bonding on the one hand and the decreased possible throughput on longer multi-hop paths on the other hand.

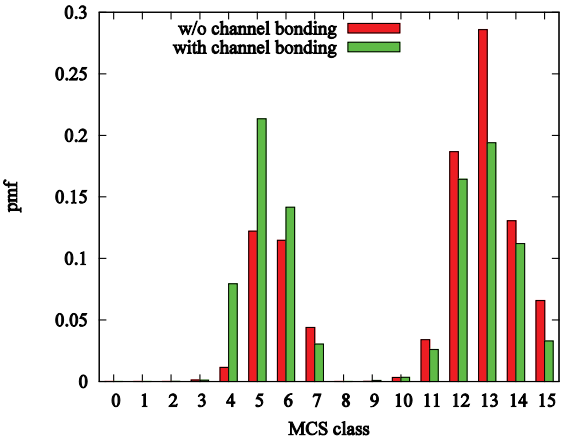

Figure 3. Fraction of utilized MCS classes in a multi-hop communication with 8 hops with and w/o channel bonding

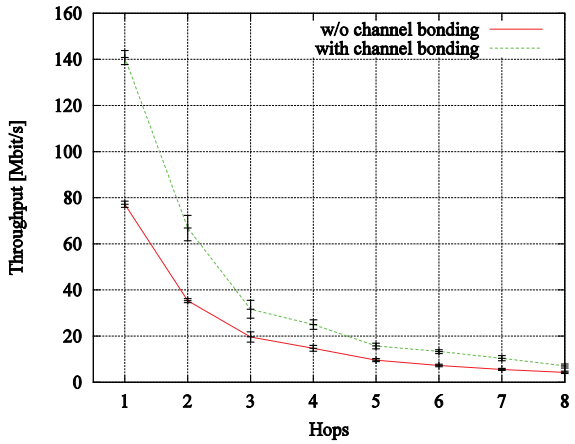

Figure 4. Throughput vs. number of hops with and w/o channel bonding without restrictions

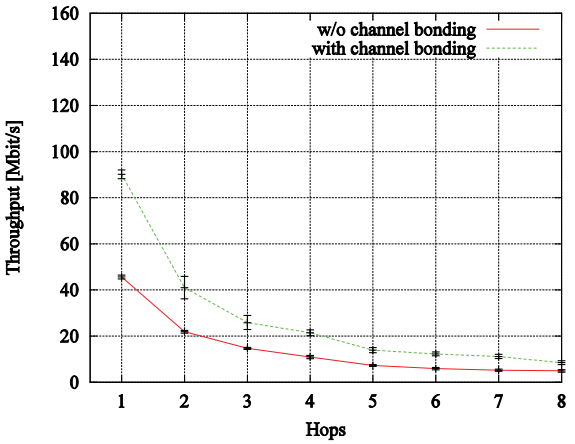

Figure 5. Throughput vs. number of hops with and w/o channel bonding restricted to single-stream mode

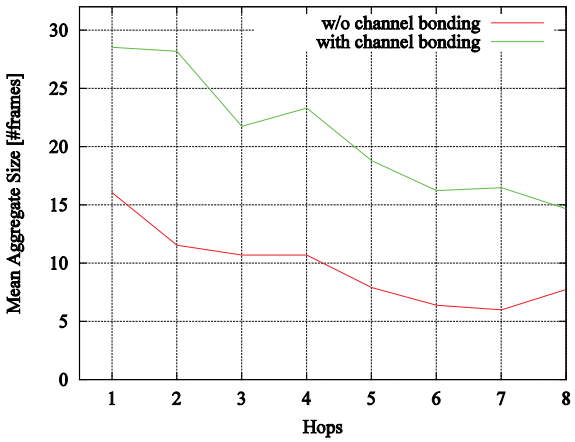

Figure 6. Mean aggregate size vs. number of hops with and w/o channel bonding

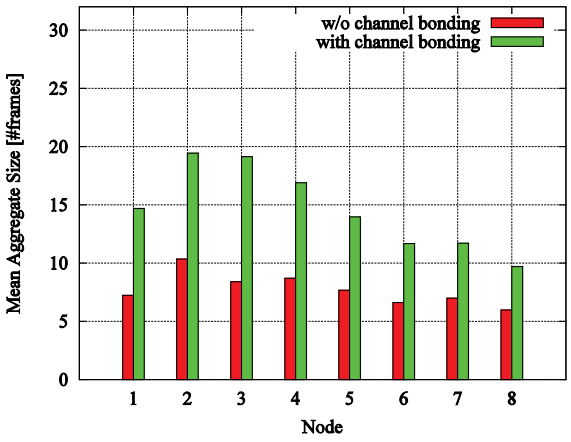

Figure 7. Mean aggregate size at each node for a multi-hop communication with 8 hops

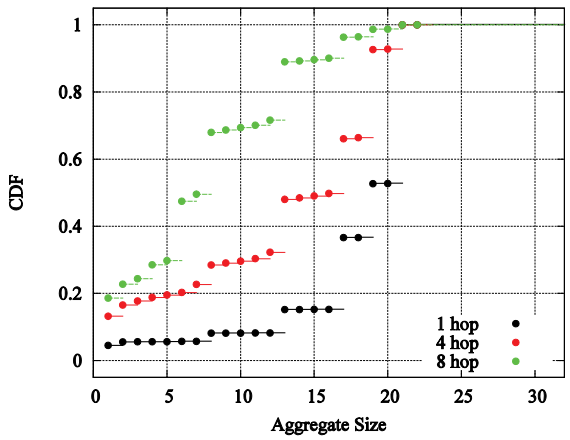

Figure 8. Cumulative distribution function of aggregate size for different path lengths and w/o channel bonding

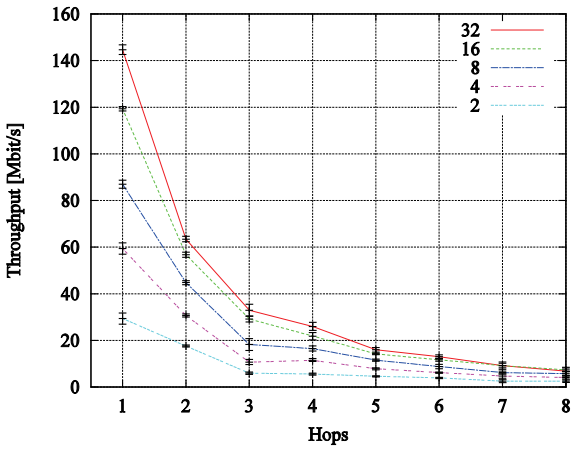

Figure 9. Throughput vs. number of hops for different maximum aggregate sizes

However, the mean aggregate size doesn't decrease as fast as the throughput on longer paths because a transmitter still can aggregate enough frames when waiting for a transmission opportunity. This is again evidence that increased medium contention is the limiting factor on longer multi-hop paths.

In Figure 7, we take a detailed look on the mean aggregate size at each node for a fixed 8-hop chain topology. We see that inside the flow the aggregate sizes differ, leading to higher aggregate sizes near the source and lower ones near the destination. We think this is because the queue saturation level on each transmitter decreases with each hop, so the last transmitter has fewer frames to aggregate.

Figure 8 plots the cumulative distribution function of the aggregation size for selected path lengths with channel bonding deactivated. We observe that while for a 1-hop flow half of all frames are transmitted in aggregates greater than 20 frames, on an 8-hop chain half of all frames are transmitted in aggregates greater than only 7 frames. Thus, the increased path length leads to much smaller aggregates and also broadens the spectrum of used aggregate sizes.

In our last experiment, we analyze the influence of the maximum aggregate size on the throughput by gradually reducing the maximum allowed number of frames per aggregate. We observe in Figure 9 that an inappropriate choice of the maximum aggregate size can potentially quarter throughput. We also observe that this effect slightly vanishes with longer path lengths. For path lengths greater than 4 hops, we observe for maximum aggregate size of 32 frames nearly the same throughput as for 16 frames. Note that for deactivated channel bonding (Figure not shown due to space limitations) this effect is even stronger. We actually observe that the throughput values for a maximum aggregate size of 32 frames are almost the same as for 16 frames for all path lengths.

V. CONCLUSION

We presented a measurement study of the multi-hop behavior of IEEE 802.11n in a real-world indoor mesh testbed for quantitatively investigating characteristics of IEEE 802.11n on multi-hop paths. In particular, we revealed details on the multi-hop behavior of the aggregation level.

The presented performance curves reveal that channel bonding nearly doubles the throughput for any fixed path length. The mean aggregate size in number of frames at each node is also doubled by channel bonding and the mean aggregate size in number of frames at each node decreases with increasing path length.

Future work investigates competitive flows and random topologies in IEEE 802.11n mesh networks to gain a comprehensive understanding of the performance of 802.11n. Furthermore, we are studying the side effects of various types of 802.11b/g background traffic.

REFERENCES

[1] S.M. Alamouti, A simple transmit diversity technique for wireless communications, *IEEE JSAC*, **16**, 1998.

[2] J. Bicket, Bit-rate Selection in Wireless Networks, Masters Thesis, MIT, Cambridge, MA, 2005.

[3] J. Bicket, D. Aguayo, S. Biswas, and R. Morris, Architecture and Evaluation of an Unplanned 802.11b Mesh Network, *Proc. ACM MOBICOM*, Cologne, Germany, 2005.

[4] S. ElRakabawy, S. Frohn, and C. Lindemann, A Scalable Dual-Radio Wireless Testbed for Emulating Mesh Networks, *Springer Wireless Networks*, **16**, 2010.

[5] S. Frohn, S. Gübner, and C. Lindemann, Analyzing the Effective Throughput in Multi-Hop IEEE 802.11n Networks, *Proc. IEEE HotMESH*, Montreal, Canada, 2010.

[6] D. Halperin, W. Hu, A. Shethy, and D. Wetherall, Predictable 802.11 Packet Delivery from Wireless Channel Measurements, *Proc. ACM SIGCOMM*, New Delhi, India, 2010.

[7] IEEE 802.11n: Standard for Wireless LAN Medium Access Control (MAC) and Physical Layer (PHY) Specifications Amendment 5: Enhancements for Higher Throughput, 2009.

[8] Iperf, the TCP/UDP Bandwidth Measurement Tool, http://dast.nlanr.net/projects/iperf/.

[9] A. Khattab, A. Sabharwal, and E. W. Knightly, Fair Randomized Antenna Allocation in Asynchronous MIMO Multi-hop Networks, *Proc. IEEE ICCCN*, St. Thomas, Virgin Islands, 2008.

[10] W. Kim, H. Wright, and S. Nettles, Improving the Performance of Multi-hop Wireless Networks Using Frame Aggregation and Broadcast for TCP ACKs, *Proc. ACM CoNEXT*, Madrid, Spain, 2008.

[11] J. Koivunen, P. Almers, V.-M. Kolmonen, J. Salmi, A. Richter, F. Tufvesson, P. Suvikunnas, A. F. Molisch, and P. Vainikainen, Dynamic multi-link indoor MIMO measurements at 5.3 GHz, *Proc. IEEE EuCAP*, Edinburgh, UK, 2007.

[12] K. LaCurts and H. Balakrishnan, Measurement and Analysis of Real-World 802.11 Mesh Networks, *Proc. ACM IMC*, Melbourne, Australia, 2010.

[13] T. Li, Q. Ni, D. Malone, D. Leith, Y. Xiao, and R. Turletti, Aggregation with Fragment Retransmission for Very High-Speed WLANs, *IEEE/ACM Transactions on Networking*, **17**, 2009.

[14] C. Papathanasiou and L. Tassiulas, Multicast Transmission over IEEE 802.11n WLAN, *Proc. IEEE ICC*, Beijing, China, 2008.

[15] I. Pefkianakis, Y. Hu, S. H. Wong, H. Yang, and S. Lu, MIMO Rate Adaptation in 802.11n Wireless Networks, *Proc. ACM MOBICOM*, Chicago, IL, 2010.

[16] K. Pelechrinis, I. Broustis, T. Salonidis, S. V. Krishnamurthy, and P. Mohapatra, Design and Deployment Considerations for High Performance MIMO Testbeds, *Proc. ACM WICON*, Maui, Hawaii, 2008.

[17] K. Pelechrinis, T. Salonidis, H. Lundgren, and N. Vaidya, Analyzing 802.11n Performance Gains, *Proc. ACM MOBICOM (poster session)*, Beijing, China, 2009.

[18] K. Pelechrinis, T. Salonidis, H. Lundgren, and N. Vaidya, Experimental characterization of 802.11n link quality at high rates, *Proc. ACM WINTECH*, Chicago, IL, 2010.

[19] D. Piazza, N. J. Kirsch, A. Forenza, R. W. Heath, and K. R. Dandekar, Design and Evaluation of a Reconfigurable Antenna Array for MIMO Systems, *IEEE Transactions on Antennas and Propagation*, **56**, 2008.

[20] V. Shrivastava, S. Rayanchu, J. Yoon, and S. Banerjee, 802.11n Under the Microscope, *Proc. ACM IMC*, Vouliagmeni, Greece, 2008.