

A Measurement Study of Multiplicative Overhead Effects in Wireless Networks

Joseph Camp, Vincenzo Mancuso, Omer Gurewitz, and Edward W. Knightly

Abstract—In this paper, we perform an extensive measurement study on a multi-tier mesh network serving 4,000 users. Such dense mesh deployments have high levels of interaction across heterogeneous wireless links. We find that this heterogeneous backhaul consisting of data-carrying (forwarding) links and non-data-carrying (non-forwarding) links creates two key effects on performance. First, we show that low-rate management and control packets can produce a disproportionately large degradation in data throughput. We define a metric for this effect called Wireless Overhead Multiplier and use it to quantify the impact of MAC and PHY mechanisms on the throughput degradation. Surprisingly, we show that these multiplicative effects are primarily driven by the non-forwarding links where, in the worst case, data packets lose physical layer capture to the overhead, yielding disproportionate throughput degradation. Finally, we show that when data flows contend in this worst-case scenario, the loss-based autorate policy is unnecessarily triggered, causing throughput imbalance and poor network utilization.

I. INTRODUCTION

Large-scale mesh network deployments are planned and underway in cities across the world. Such networks employ a multi-tier architecture consisting of an access tier that connects clients to mesh nodes, a backhaul tier that wirelessly inter-connects mesh nodes, and a capacity injection tier that provides high-performance wireless links from some backhaul nodes to fiber gateways. Combining research and community access objectives, we have deployed and operate a multi-tier mesh access network, Technology For All (TFA), that serves 4,000 users in a densely populated, urban neighborhood.

In the deployment of a multi-tier mesh network, primary consideration is given to the performance of the forwarding links, i.e., links selected by the routing protocol to forward traffic to and from wired gateways. For example, prior work has studied the tradeoff between node spacing and the performance of the resulting links and multi-hop paths [1]. However, since the set of nodes within the main forwarding path use a shared medium, the addition of mesh nodes along the forwarding path also creates a large number of *non-forwarding* links, or links that are not selected or cannot be selected by the routing protocol to forward data. In the strictest sense, every node forms a link with every other node, even if the resulting link yields near negligible interference. In any case, the resulting connectivity matrix of forwarding




Fig. 1. Heterogeneous links inherently exist in mesh networks even in an idealized topology and propagation environment.

and non-forwarding links within a mesh topology is vastly heterogeneous in quality due to relative differences in spacing and wireless propagation characteristics among nodes. We term this matrix *heterogeneous backhaul connectivity*.

The aforementioned heterogeneity is unavoidable, whether these non-forwarding links are foreseen or not during the design process. For example, consider a hexagonal topology in which all neighbors are one-hop and have identical distance to the gateway, as shown in Fig. 1. Even if the propagation environment is homogeneous (e.g., with a uniform path loss), the links formed within the topology are not identical. Although the forwarding links to the gateway could in principle be homogeneous, the non-forwarding links are inherently *heterogeneous* due to the geometry. Thus, an effective mesh topology (non-triangle) necessarily yields heterogeneous links. In a real-world deployment, even greater link heterogeneity exists due to topological irregularities, non-uniform path loss, and the presence of shadowing within the environment.

In this paper, we experimentally explore the interaction of heterogeneous backhaul connectivity with both overhead and data traffic. Our contributions are two-fold.

We first consider control traffic employed by network management functions such as routing, client association, and link establishment. We refer to such traffic as “overhead” and explore the effect of injected overhead on the throughput of a data flow generated by a different node. Unfortunately, we find that overhead can have a substantially more problematic effect than merely subtracting an equivalent throughput from the data flow. We introduce the Wireless Overhead Multiplier (WOM) to characterize this effect and experimentally show that WOM is controlled primarily by the non-forwarding links, secondarily by the relative strength of the forwarding links, and eventually by the protocol set.

Our experiments yield a worst-case scenario in which the overhead traffic consistently wins a physical layer capture over the data traffic when the data-sender and overhead-injector are mutually out of range. In this case, injection of merely 10

J. Camp and E. Knightly are with Rice University, Houston, TX (<http://networks.rice.edu>). V. Mancuso is with DIEET, Università di Palermo, Italy. O. Gurewitz is with Department of Communication Systems Engineering, Ben Gurion University of the Negev, Beer Sheva, Israel. All four authors were with Rice while performing the research for this work.

This research was supported by NSF grants CNS-0331620 and CNS-0325971 and by the Cisco Collaborative Research Initiative.

kbps of overhead traffic decreases data throughput by over 500 kbps, yielding a multiplicative effect of over 50 times the overhead rate. Thus, while minimizing overhead traffic is a well accepted tenet of protocol design, our results establish a severe reduction in data-path performance with even moderate overhead. We experimentally investigate overhead effects on a single data flow by characterizing the impact of a single source of overhead as a function of its location. In addition, we show that the effect of multiple overhead-injectors can be analyzed by means of the results obtained for the single injector.

Finally, we explore the impact of backhaul connectivity on contending data flows. We show that in some scenarios, WOM-like effects also exist for data traffic. However, the high rate and large packet sizes of data as compared to overhead yield a significantly different interaction, namely the compound effect of physical capture and adaptation of the modulation rate via Auto Rate Fallback [2]. Two relevant aspects of ARF have been previously studied. First, because ARF uses acknowledgments to infer that a link quality is poor, it can suffer from incorrectly interpreting collisions due to other reasons (e.g., hidden terminals) as being due to channel fading. Our experiments indicate that this effect indeed is prominent in practice. Second, because ARF is “packet fair” vs. “time-share fair” [3] it suffers from a “performance anomaly” in which the lowest rate link controls all flows’ throughputs [4]. Our experiments indicate that, in practice, this effect is overwhelmed by the capture effect. In other words, high-quality links can capture the channel to overwhelm low-quality links and the performance anomaly is not observed.¹

II. TECHNOLOGY FOR ALL MESH ACCESS NETWORK

In this section, we present the TFA network and our experimental set-up and methodology.

A. Network Description

The TFA network is a multi-tier mesh access network deployed in a densely populated, single-family residential neighborhood. Currently, 18 backhaul nodes are predominantly deployed on single-story residences with the exception of three schools, two businesses, and a public library. The spatial distribution of the backhaul nodes are shown in Fig. 2 and are graphically connected if a wireless link can be established between two nodes. All of the wireless links are omni-directional in nature with the exception of three long-haul directional links, pictured darker. The backhaul nodes share Internet bandwidth from a single 100 Mbps fiber and currently serve 4,000 users. The coverage area is 3 km² and has a population density of 4,760 residents per km². For details of the hardware and community, refer to [1] and <http://tfa.rice.edu>.

B. Experimental Set-up

In our experiments, we use the existing TFA backhaul nodes deployed at locations pictured in Fig. 2. We isolate a newly

¹The term “capture” has been used in the literature to refer both to MAC-layer capture [5] and physical-layer capture [6]. In this paper, we refer to capture only in reference to the latter effect.




Fig. 2. Connectivity graph of the TFA backhaul with appropriate scaling for distance between nodes. There are 4,000 residential users (not shown).

formed experimental mesh topology by switching the channels of the gateway and the desired nodes under test to form an experimental “submesh.” We use this submesh to measure the effect of adding backhaul overhead-injectors and data-senders to the network. Directional long-haul links serve as virtual gateways, providing Internet to the remainder of the mesh topology, i.e., backhaul nodes that are not involved in the experiment. We perform over 200 experiment trials on various submesh topologies at off-peak times during the night (3am–6am) to minimally obstruct real user traffic and conversely, so that user traffic has a minimal effect on our measurement results. We use *Kismet* and *tcpdump* to collect MAC-level traces at selected network nodes. Unless otherwise specified, the RTS/CTS mechanism is disabled and autorate is enabled, which represents our default operational setting.

III. MULTIPLICATIVE OVERHEAD EFFECTS OF HETEROGENEOUS BACKHAUL CONNECTIVITY

Control messages inherently must be exchanged for network management protocols such as routing, client association, and backhaul link maintenance. We define all such messages as “overhead” and omit per-data-packet overhead such as packet headers, RTS/CTS messages, etc. Overhead can be generated from devices within a network (controllable by the network operator) or from external devices such as residential APs and clients (uncontrollable by the network operator).

While a well-understood tenet of protocol design is to restrict overhead traffic to a minimum, typically via use of low-rate periodic or on-demand small-sized messages, in this section, we show that despite having low rate, overhead can profoundly degrade network performance. Specifically, an overhead rate of λ can reduce the data throughput on a nearby link by up to 50 times λ . Here, we measure an initial scenario, characterize the overhead, and present our measurement methodology to show the factors driving such effects.

A. Diverse Overhead Effects

In this section, we quantify the impact of overhead traffic on data throughput. To achieve this, we design an experiment in which we compare the throughput of a single link with and without the overhead induced by the surrounding nodes. In the experiment, we measure the *achievable throughput* defined

as follows. Given a network N , and a sender-receiver pair $s, r \in N$, consider a fully backlogged flow $f_{s \rightarrow r}$ from node s to node r . The *achievable throughput* of the flow $f_{s \rightarrow r}$ is the throughput $t_{s \rightarrow r}^N$ achieved when all nodes in $N \setminus \{r, s\}$ only transmit overhead. Achievable throughput is defined for a particular protocol set (e.g., long-lived UDP flows with 1500 byte packets, over 802.11 with no RTS/CTS, and autorate enabled).

To eliminate known throughput degradation effects such as [1], [7], we first measure only single-active, one-hop flows where the user activity of the system is negligible. Furthermore, we concentrate on the effects of overhead only on high quality links (i.e., links that can send at the highest modulation rate). Thus, we measure the throughput degradation of each link's achievable throughput due to the injected backhaul overhead traffic.

In particular, we select a single one-hop backhaul node near the gateway (see Fig. 2) to send backlogged UDP traffic when all surrounding nodes are disabled and measure the UDP achievable throughput. We then measure the achievable throughput of the same sender-receiver pair in the presence of overhead from surrounding nodes, i.e., neighboring nodes in the network are enabled but allowed to transmit only control traffic. For both measurements, we have identical hardware configurations for all nodes (200 mw transmit power, RTS disabled, autorate enabled), and hold the traffic type constant (1500 byte, constant bit rate, UDP traffic). We repeat the three-node experiment sequentially for each node that is one hop from the gateway.




Fig. 3. Achieved throughput with and without overhead (isolated) injected from the TFA network.

Fig. 3 shows the throughput degradation that each node experiences where the x-axis is the backhaul node and the y-axis is the achievable throughput. For each node, the left bar represents the achievable throughput in isolation (when no other nodes are transmitting overhead) and the right bar represents the achievable throughput when the network overhead is injected.

As can be seen in Fig. 3, there are two dramatic effects from the overhead. First, there is a sizable portion of each achievable throughput in isolation that is lost on each link (at least 1/5 of the throughput is lost on each link). Second, the throughput degradation is vastly different among links caused by the presence of overhead within the network. Specifically, the throughput degradation ranges from 850 kbps in the best-case to over 1700 kbps in the worst-case. Since the only difference between the setup of the measurements taken for the two bars for each node is the presence of TFA network

overhead, the throughput degradation must be associated with the overhead injected by TFA. In order to verify these results, we repeated the same experiment on all channels. Indeed, in all channels, we observed the same trend which verifies the cause of the effect is the network overhead and eliminates the possibility that the two effects perceived in Fig. 3 are exclusive to the channel used by TFA or due to interaction with external networks operating on the same channel.

As previously explained, the throughput degradation is solely related to the overhead injected by the nodes of the TFA network. Furthermore, in all tested pairs the receiver is the gateway which sees the same number of transmitters (overhead-injectors) and the same environment (other noise-injectors, etc.) across all measurements. Also, since the hardware platform of all senders tested is identical (transmission power, autorate policy, RTS/CTS mechanism, etc.), the differences in throughput degradation caused by the overhead must be due to the location of each transmitter, i.e., topological differences seen by each transmitter. More specifically, the throughput degradation experienced by each sender is correlated to the quality of the bidirectional links that are formed between the sender and the other nodes (receiver and overhead-injectors).

In the remainder of the paper, we investigate this correlation between the throughput degradation caused by network overhead and link heterogeneity.

B. Heterogeneous Non-Forwarding Links

The only difference driving the heterogeneity in overhead effect is the varying spatial location of overhead-injecting nodes to the data-sender. These links between transmitters which are not intended to communicate directly (non-forwarding links) are inherent within the topology (i.e., not planned within the design of the forwarding links of the network). Thus, these non-forwarding links vary greatly in quality compared to the data-carrying or forwarding links. Such non-forwarding links impact the data transmission whether causing the node to defer at the transmitter or yielding simultaneous transmissions resulting in collisions or capture effects. For example, node n7 can cause node n4 to defer since the two nodes are able to decode one another's packets. On the other hand, n4 and n8 are unable to decode each other's packets or even sense each other on the medium and hence collide.

Thus, the difference in overhead effects is caused by the differing nature of these links between sender and non-receiver neighbors. We define the resulting connectivity matrix of vastly heterogeneous non-forwarding and forwarding links within a mesh topology as the *heterogeneous backhaul connectivity*. We now define a term to quantify the multiplicative overhead effects caused by the heterogeneous backhaul connectivity.

C. Wireless Overhead Multiplier

To formally define WOM, consider a sender-receiver pair s, r and a set O consisting of nodes which are primary interferers to s and r , as depicted in Fig.4. Let λ_O denote the cumulative mean rate of the overhead transmitted by the




Fig. 4. Node s sends data and overhead (OH) to node r while all other nodes within set O transmit only OH.

nodes in O . Of course, $t_{s \rightarrow r}^N$ is related to the set of active interfering nodes O , i.e., the more interfering nodes are active, hence transmitting overhead, the lower the expected $t_{s \rightarrow r}^N$. Accordingly, we evaluate the impact of the overhead due to O on flow $f_{s \rightarrow r}$ by comparing the achievable throughput $t_{s \rightarrow r}^N$ in isolation (i.e., $N = \{s, r\}$ and nodes in O are not active), with the value of $t_{s \rightarrow r}^N$ when the interferers are active (i.e., $N = \{s, r\} \cup O$):

$$W_{s \rightarrow r} = \frac{t_{s \rightarrow r}^{\{s, r\}} - t_{s \rightarrow r}^{\{s, r\} \cup O}}{\lambda_O} \quad (1)$$

Eq. (1) gives a measure of the achievable throughput degradation normalized to the injected overhead. Notice that, since overhead can be due to pure MAC frames (e.g., beacons) as well as to IP packets (e.g., routing messages), we include in λ only the MAC throughput, i.e., we take into account only the payload of overhead frames. Hence, ideally, the protocol overhead causes a degradation of the achievable throughput equivalent to the air-time utilization of overhead traffic, which is greater than λ_O . For example, a short unicast (90-byte) IP message sent at maximum modulation rate (11 Mbps lasting $58 \mu s$) incurs a per-packet overhead of a preamble (at 2 Mbps lasting $192 \mu s$) and a 30-byte MAC header (at 11 Mbps lasting $22 \mu s$) plus the 14-byte ACK ($192 \mu s$ for the PHY preamble, and $10 \mu s$ for the ACK MAC frame at 11 Mbps). Hence, also considering the spacing between frames (SIFS and DIFS), a 90-byte packet flow uses a gross bandwidth of 11 Mbps to carry 1.3 Mbps, i.e., the overhead consumes $11/1.3$ times its nominal bandwidth λ_O . Analogously, a 1500-byte payload transmitted at 11 Mbps yields an average transmission rate of approximately 7.9 Mbps, i.e., an actual overhead rate of $1.4 \cdot \lambda_O$. Thus, the ideal expected WOM value caused by acknowledged frames ranges from 1.4 to 8.5, depending on the size of the overhead payload.

However, we find that the WOM value can range from near 0 to over 50. Further, we show that the WOM is controlled by effects due to the heterogeneity of the quality of all links formed between nodes s, r , and the interferers in O . In order to understand the basic interaction of links, we first investigate the WOM effect within a topology of three nodes: the data-sender (s), the data-receiver (r), and the overhead-injector (o). Thereafter, we show the compounding effects of more complex topologies.

IV. ISOLATING LINK EFFECTS

In this section, we explore the effect of a single overhead-injecting node as a function of the link quality to a data-sender and a data-receiver. Following the 802.11 standard, nodes behave differently according to differing link qualities with respect to other transmitters. Thus, we classify links according to the transmitter behavior specified in IEEE 802.11 and isolate the overhead effects due to different node behaviors using a three-node topology as shown in Fig. 5. Node s represents the data-sender, node r represents the data receiver, and node o represents the overhead-injecting node. Links s, r and o, r both are able to achieve transmissions at the maximum rate, while link o, s can vary.




Fig. 5. Node s sends data and overhead (OH) to node r while node o transmits only OH.

A. IEEE 802.11 Node Behaviors

The standard describes three different behaviors within for medium access: (i) if a node is able to decode a transmission of another sender, it NAVs (according to the physical or virtual carrier sense mechanism), (ii) if a node is able to detect channel activity, it defers until the channel is free and additionally defers its transmission for an Extended Inter Frame Space (EIFS) period which covers the longest possible ACK duration, and (iii) if a node is unable to detect channel activity, it transmits according to the normal backoff mechanism. Correspondingly, we classify each node pair (i.e., the link between the two nodes) according to their degree of connectivity as defined by the standard: (i) transmission range, (ii) carrier sense range, and (iii) out of range. While variation in channel quality can cause links to change their class over time, each individual packet is within a single class according to the MAC behavior.²

IEEE 802.11 Off-the-Shelf Card Behavior. We begin our investigation by testing the off-the-shelf hardware for the node behavior described in the 802.11 standard to enable detection of a particular TFA link class. Determining that nodes are in transmission range can be achieved simply by ensuring that beacons (sent at the base rate) are successfully received. However, distinguishing between carrier sense range and out of range classes requires experimentation since the MAC state machine is not directly observable. Thus, we next design an experiment to distinguish between these two classes.

If s adopts the energy detection behavior described in the standard, when the energy level is above a given threshold, it will defer transmission via physical carrier sensing. To experimentally find the energy detection threshold, we use the

²For a statistical description of the links in TFA, refer to [8].

configuration depicted in Fig. 6 where the sender-receiver pair s, r of the data flow communicate over the air, and the noise generator n is hard-wired to the sender.




Fig. 6. 802.11-behavior experiment set-up for the off-the-shelf wireless cards.

In the experiment, we send a UDP flow from s to r at a constant physical layer rate of 2 Mbps. The link from s, r is held constant. We observe the behavior of the achievable throughput of s, r as a function of the noise level generated by n . Thus, any change in the throughput at r is caused by the behavior of s , i.e., if the noise is above an energy detection threshold, s defers, driving the achievable throughput to 0.




Fig. 7. 802.11 card behavior when noise is injected at the transmitter only.

Non-Existence of Carrier Sense Range. Fig. 7 depicts the throughput of the data flow from s to r where the x-axis is the level of generated noise. The noise source is a modulated sine wave within the spectrum of the 802.11 channel used in the experiment. We observe a dip from 1230 kbps (the achievable throughput when the noise source is disabled) to 360 kbps at -60 dBm. We observe with *Kismet* that the throughput decrease is due to the deafness of the transmitter to hear the ACK, leading to excessive backoffs and retries of the same application layer packet. Thus, there is no energy detection threshold.

Therefore, we find that the chipset/driver used in TFA (Prism/HostAP) defers only when another packet in the air is able to be decoded and will not defer due to noise alone. Note that this behavior is compliant with one of the modes available for CCA procedure described in the IEEE 802.11 standard [9], i.e., carrier sense without energy detection. Furthermore, this is a common choice for vendors, e.g., another well-known chipset/driver, the Atheros/MadWiFi, operates in the same way. Hence, in TFA there are only two link classes.

B. WOM in Three-Node Topologies within TFA

In this section, we experimentally study the WOM effect defined by Eq. (1) on the TFA network in accordance with the TFA link classes. To achieve this, we systematically isolate three node topologies from the same nodes involved with the experiment shown in Fig. 3, all other nodes are disabled.

Specifically, we perform extensive measurements to form a data set from these three-node topologies consisting of both TCP and UDP data traffic of 1500 bytes from the data-sender s

to the data-receiver r . As observed via *tcpdump* and *kismet*, the overhead traffic sent from the third node o consists of 90-byte packets (on average) at approximately $\lambda = 10$ kbps. More than 90% of the overhead traffic consists of unicast frames, and user traffic is negligible. AutoRate Fallback is enabled in all experiments unless otherwise specified.




Fig. 8. WOM considering the link class (transmission range or out of range) of the link between the WOM-inducing node o and the data-sender s .

Fig. 8 shows the average WOM values with error bars representing one standard deviation above the value for our data set where nodes are within transmission range and out of range with TCP data traffic. We find that the average WOM induced by an overhead-injecting node within transmission range is 4.6. Further, we observe that the nodes out of range exhibit a much larger average (11.8) and variance in the WOM values as compared to the transmission range.

The results in the transmission range case are not surprising. In fact, due to the perfect coordination between the nodes (beyond the negligible propagation delay), the WOM value is in the range predicted for the ideal case discussed in subsection III-C. In contrast, when the data-sender and the overhead-injecting nodes are out of range, there is a lack of coordination that yields significantly larger WOM values as we now explore.

C. WOM in Out of Range Class

Within the out of range class, simultaneous transmissions occur causing various effects: (i) collisions resulting in loss, (ii) retransmissions, and (iii) the physical layer capture effect. We now describe these effects within the context of the MAC (i and ii) and PHY (iii) layers.

MAC Effects. In a CSMA MAC, simultaneous transmissions can collide at a mutual receiver, resulting in loss and retransmissions. Since the optional RTS mechanism was designed to avoid such collisions, we investigate the WOM effects with and without this collision avoidance mechanism.

With RTS disabled, the cost of a single retransmission is approximately one doubled backoff period plus the packet period. In the case of low-rate overhead, each data packet from s is unlikely to collide with more than one overhead packet. Hence, the collision rate of the system is approximately equivalent to the overhead packet injection rate. For example, to retransmit a 1500 byte packet at 11 Mbps, it takes approximately 2.2 ms, on average, including DIFS, SIFS, ACK, backoff and PHY overhead. Hence, an overhead of 90 byte packets at 10 packets per second (i.e., λ equals 7.2 kbps) reduces the rate of successful transmissions of s to r and yields a WOM value of over 20 for UDP traffic. However, because

Fig. 8 indicates substantial variation from this point, hidden terminal collision effects alone are insufficient to characterize the WOM value.

With RTS enabled, the cost of collision is reduced to one doubled backoff period plus the RTS/CTS exchange duration. We now compare the aforementioned cost of collision to the cost of the additional signaling imposed by the use of the RTS/CTS mechanism to the gains of the reduced cost. To compare this, we show the case where two nodes (n4 and n8) are out of range. We use n4 as the data-sender s and n8 as overhead-injector o . We measure the induced WOM with and without the use of the RTS/CTS mechanism with TCP data traffic.




Fig. 9. WOM (left) and aggregate TCP (right) considering use of RTS mechanism in an out of range scenario.

Fig. 9 (left) depicts measurements of WOM over multiple trials for node n4, with an out of range overhead-injector, n8. Indeed, the WOM is reduced by the use of the RTS/CTS mechanism. However, note that the protocol set for a given WOM (see the definition in III-C) has changed, thereby altering the achievable throughput used for reference flow. Since the RTS/CTS mechanism induces *per-packet* overhead, the use of the protocol set here with TCP traffic with RTS enabled has lower achievable throughput than TCP traffic with RTS disabled. The induced per-packet overhead of RTS used to reduce the WOM produces a net loss of aggregate throughput. Namely, the achievable throughput of n4 is 2.5 Mbps with RTS enabled and 3.3 with RTS disabled, *after* WOM is taken into account. In summary, our measurements indicate that while use of RTS/CTS reduces WOM, its increased per-packet overhead yields a net throughput reduction for data traffic.

Joint PHY/MAC Effects. Throughput and MAC behavior are strongly influenced by physical layer capture [6]. Thus, we next establish the existence of capture in the TFA network and explore its impact on WOM.

First, since it has been shown that ARF causes throughput imbalances in the hidden terminal scenario [10], we fix the physical layer rate of the transmission to the base rate (2 Mbps) to eliminate these effects. Next, we measure the achievable throughput of each one-hop backhaul node s from the gateway r in isolation and in the presence of one out of range overhead-injector o . We also record the differences in SNR at the gateway between the two transmitters.

Fig. 10 shows the WOM value for each of the differences in SNR where a positive value indicates s has a more powerful SNR at r than o . The results indicate a bimodal relationship in the WOM values for the positive and negative SNR differences. More precisely, when the SNR difference is positive, the WOM value is approximately 1, indicating that




Fig. 10. WOM of out of range links considering relative RSSI at data-receiver r from data-sender s and WOM-inducing node o .

the overhead losses experienced by the data sender are less than the actual injected overhead. However, when the SNR difference is negative, the WOM value ranges from 6 to 12. We conclude that capture effect occurs with a difference in SNR of greater than 0.

To show that this bimodal behavior is due to the capture effect, we now evaluate the WOM associated with two specific out of range nodes. We use the first node n7 as a data-sender, and the second node n2 as an overhead-injector and measure the WOM. The SNR from n7 is 3 dB greater than from n2 at GW. We then repeat the experiment after switching the roles of the nodes.




Fig. 11. Asymmetry of WOM of two nodes with respect to one another.

Fig. 11 shows the WOM for the two experiments for both TCP and UDP fully-backlogged traffic and physical layer rate of 2 Mbps. Node n2 has a WOM value of 9.2 and 7.6 for UDP and TCP, respectively; while n7 has a WOM value of 0.9 and 0.6, respectively. Hence, we find that the severe asymmetry exists across both traffic types. Regardless of the traffic type, the out of range class must be split into two subclasses to characterize the WOM behavior.

D. Discussion

We now post-process Fig. 8 considering a positive SNR difference (or capture win) or negative SNR difference (or capture lose) by the data-sender s . Fig. 12 illustrates the net effect of capture and depicts the WOM values for the two cases as to whether the data transmitter s wins or loses the capture. The figure indicates that despite node o being in the same out-of-range class, the WOM value can be as small as 1 (capture win) or as large as 25 (capture lose). The observed physical layer capture effect on WOM explains the asymmetry shown in both Fig. 9 and Fig. 11.

In summary, the primary factor that controls the aforementioned WOM classes/subclasses is the non-forwarding links, i.e., the level of coordination the data-sender has with the overhead-injector. If the two transmitters are out of range, the




Fig. 12. WOM considering the different effects of physical layer capture effect within the out of range link class.

secondary factor is the relative quality of the forwarding links, i.e., the relative SNR at the mutual receiver which drives the capture effect. From these two factors, clear WOM modes can be established. Finally, the behavior within the modes is driven by a tertiary effect, the protocol set, consisting of the traffic type (e.g., TCP or UDP) and protocol parameters (e.g., the usage of RTS/CTS).

For example, we can reconsider the RTS results presented in Fig. 9 according to the discussion above: The RTS/CTS mechanism (tertiary effect) is unable to completely reduce WOM to the values associated with the transmission range scenario since the RTS messages are also captured (secondary effect) at the mutual receiver, thereby reducing the ability of the collision avoidance mechanism to counter the hidden terminal problem (primary effect).

Finally, although we cannot show the carrier sense WOM behavior within the TFA hardware and environment, the expected values are similar to the transmission range class as verified by ns-2 simulation.

V. CUMULATIVE LINK EFFECTS

In this section, we use the findings from the previous section to explain more complex scenarios. We first characterize injected overhead as it scales with the number of TFA backhaul nodes. We then measure the achievable throughput with an increasing number of these overhead-injectors.

Scaling Overhead Transmitters. Before focusing on the effect of the injected overhead, we must first understand the actual overhead that is being transmitted by a particular overhead-injector. To achieve this, we passively sniff the channel using *Kismet* on a wireless node located next to the gateway with no backhaul nodes enabled. We then sequentially enable first the gateway and then each one-hop node around the gateway. For each backhaul node that is enabled, we allow the routing protocol (AODV³) to reach steady-state before enabling the next backhaul node.

We observe that 99% of the overhead belongs to the category of periodic low-rate single-hop messages, as link failures and other events requiring flooding occur rarely. Fig. 13 depicts the overhead messages (AODV and Beacon frames) in steady-state for the number of backhaul nodes depicted on the x-axis. As can be seen in the figure, each node adds approximately 10 kbps overhead to the network (i.e., the




Fig. 13. Overhead observed at GW as the overhead-injectors scale.

overhead grows linearly with the number of nodes). Note that due to the presence of other WLAN networks within the same TFA environment, we observe an overhead floor of approximately 20 kbps (labeled “ref” in Fig. 13).

Scaling WOM Effects. Finally, we study how a mesh network’s overhead effects scale as overhead-injecting nodes are sequentially enabled. To achieve this, we measure a data sender-receiver pair $n4, gw$ and sequentially enable nodes to the topology until all the one-hop nodes from Fig. 3 are injecting overhead.




Fig. 14. Throughput degradation at each step when sequentially adding overhead per node within the topology as experienced at $n4$.

Fig. 14 shows the throughput degradation experienced by the TCP data flow from $n4$ to GW when the specified backhaul node on the x-axis is enabled. Specifically, the throughput degradation corresponds to the achievable data throughput before and after a single overhead-injector is sequentially enabled. We see that $n4$ has a wide span of throughput degradation values, from 20 to 520 kbps. The lowest value, 20 kbps, is experienced when the out of range overhead-injector, $n2$, loses the capture with respect to the data flow and thus, $n4$ transmits data without knowledge of the overhead from $n2$. In contrast, the highest value, 520 kbps, is experienced when the out of range overhead-injector, $n7$, wins the capture with respect to the data flow and thus, $n4$ has excessive timeouts and backoffs for data packets causing wasted air-time and losses. While $n6$ and $n8$ are out of range, it is unclear whether they would win or lose a capture since the SNR at the gateway is less than 1 dB different for transmissions from $n4$, $n6$, and $n8$. Thus, the values of 90 and 130 kbps, respectively, are between the extremes of capture win and lose. Lastly, nodes $n1$ and $n3$ are within transmission range of $n4$ and correspondingly result in throughput degradations of 40 and 80 kbps, respectively. Thus, we find that as network size scales, injected overhead increases linearly, yielding cumulative degradation in data throughput in accordance with the WOM link relationships established in the previous section.

³The IEEE 802.11s standard uses a variant of AODV in which the control traffic is unchanged yet has a different link metric. Hence, its basic overhead structure is the same as reported here.

VI. HIGH-RATE FLOW EFFECTS OF HETEROGENEOUS BACKHAUL CONNECTIVITY

In this section, we study the effect of heterogeneous backhaul connectivity on mutually contending data flows. Similar to the WOM scenario, each data flow experiences the presence of other flows as if high-rate overhead is injected in the network. Here, the key difference is that the size and spacing of the packets are sufficient to cause consecutive losses, triggering the autorate policy, ARF.

We find that in the same scenario as for the worst-case WOM (out-of-range senders with physical layer capture present) that ARF lowers the transmission rate, extending the time of transmission for the capture-losing node, thereby, lowering the probability of that node fitting the packet in idle periods. Thus, the compounding effect of the autorate policy and capture effect amplifies the throughput degradation beyond the mere combination of results presented on capture effect [6] and ARF [10]. More details on the impact of link classes on the multi-flow interaction can be found in [8], omitted here for lack of space.

Autorate Penalty due to Capture. When two senders are out of range and transmit to the same receiver, the data-sender winning the capture dominates the channel and the losing sender can experience extreme starvation [6]. Yet, it is known that collisions and packet loss cause false positives within loss-based autorate mechanisms, misinterpreting the channel state [10]. Thus, ARF can drive the loser of the capture effect to decrease its transmission rate, and in turn to extend the duration of its transmissions, which makes the capture-loser's frames even more prone to collision.

To verify this hypothesis, we design an experiment in which we select two data-senders, s_1 and s_2 , such that the former wins the capture at r . We initiate simultaneous TCP flows from both nodes to r . We first run the experiment with the autorate mechanism enabled, and we then fix the transmission rate to 11 Mbps and repeat the experiment.




Fig. 15. Throughput in contention scenario with two out of range senders out of range from one another for the two cases: autorate enabled (left), and 11 Mbps fixed rate (right).

Fig. 15 (left) demonstrates the case in which the autorate is enabled. The individual and aggregate throughput are reported in the figure, as well as the achievable throughput of each flow. Note that the probability of success for the capture losing node is close to zero. In fact, since RTS/CTS is off, in order to succeed, a packet transmitted by s_2 (capture loser) has to fit in the time interval between two packets of the capture winner node, i.e., a few mini-slots. Accordingly, Fig. 15 (left) depicts that flow $s_2 \rightarrow r$ starves and receives less than 1% of the aggregate throughput.

We now disable the autorate mechanism in the same experimental set-up with two out of range senders and fix the transmission range to the physical rate the TFA links were designed for, i.e., 11 Mbps. Fig. 15 (right) depicts the results for this case. We find that the aggregate throughput of the multiplexed flows is the same as the throughput for an individual flow, although the weaker transmitter now achieves 1.0 Mbps of throughput.

Compounding Effects with System Scale. To explore the impact of ARF choosing the physical layer rates poorly, we now increase the number of contending flows. We activate up to six long-lived TCP data flows from nodes which are one-hop from the TFA gateway r , and all flows are directed to r . Flows are initiated sequentially, according to the RSSI of the transmitter received by r . Since the capture effect is present, we proceed from the weakest sender to the strongest. The flows are initiated at 3 minute intervals until all six flows are active.




Fig. 16. Six flows sequentially added in contention with ARF enabled.




Fig. 17. Six flows sequentially added in contention with a fixed physical layer rate of 11 Mbps.

Figures 16 and 17 depict the results of the experiments and show the aggregate throughput of all flows at each step of the experiment for ARF enabled and for a fixed physical layer rate of 11 Mbps, respectively. The throughput of each flow is represented as segments of each bar within the figure. The results indicate that for the case of the maximum number of contenders, the ARF mechanism degrades aggregate throughput by more than 50% compared to ARF disabled and a fixed physical layer rate of 11 Mbps.

We conclude that the ARF mechanism is unneeded in a network where forwarding links are well designed and the presence of heterogeneous non-forwarding links can only drive loss-based ARF to erroneous decisions.

VII. RELATED WORK

Scaling Control Overhead. [11] showed that on-demand routing protocols (e.g., AODV) scales well in ad hoc networks, namely, that there is a linear induced overhead when adding nodes, yet it is susceptible to failing to meet latency and

QoS requirements. While our measurements confirm the linear scaling of AODV and other overhead sources (e.g., beacons) in mesh networks, our results also unveil an unexpected multiplicative effect on system capacity losses that greatly exceed the AODV overhead due to link heterogeneity.

Physical Layer Capture Effect. [6], [12], and [13] showed the presence of physical layer capture with indoor measurements. Further, [13] showed that the RTS/CTS handshake is unable to prevent unfairness in the form of channel capture. Here, we show that inherent heterogeneity within mesh backhaul connectivities produce physical layer capture resulting in long-term throughput imbalance of backhaul sharing.

Auto Rate. In [4], it was shown that the performance anomaly of adding low-quality links within a WLAN severely decreases the throughput of high-quality links. According to [4], the poor quality links utilize the channel most of the time and degrade the throughput of the high quality links due to fair packet sharing rather than fair time sharing [3]. We experimentally show that due to the physical layer capture effect, packet fairness is not achieved, and a severe throughput reduction is due to the ARF mechanism within heterogeneous backhaul connectivity. In [14], it was shown that an autorate policy is potentially detrimental to the performance of multihop paths, depending on the routing protocol. We further show that the autorate mechanism can force losses in aggregate throughput of the system even in single-hop contention scenarios due to the incorrect choosing of appropriate rates.

Addition of Mesh Network Nodes. [15] concluded that the addition of mesh nodes randomly distributed within a given coverage region is shown to increase throughput and connectivity. However, [15] they do not report protocol overhead nor measure its effects. Our work shows that increasing the mesh node density (and thus, number of overhead-injectors) can lead to severe aggregate throughput reduction. Further, [1] studied node placement and the tradeoff between spacing of nodes, link quality, and traffic matrices using chain topologies, but the effect of non-forwarding links was not explored. In particular, we highlight the induced losses caused by forwarding and non-forwarding links by non-data-injecting nodes as well as fully backlogged nodes.

Measurement Studies. Measurement studies have been performed on single-tier (i.e., wireless backhaul only) mesh networks considering appropriate routing metrics [16], [17], link capture effects [18], and analysis of different protocol implementations (802.11b and 802.11g) [19]. Likewise, prior work has shown that as the density of mesh nodes increases, the forwarding path can have higher quality links, thus increasing connectivity and throughput of the network at-large [15]. Other measurement studies have been performed on large-scale campus WLANs [20] which evaluated traffic load and user behavior. In contrast, we focus on the relationships of forwarding and non-forwarding links across mesh backhaul networks and the effect of induced load on such a backhaul.

Thus, no prior work has considered the effects caused by the heterogeneous backhaul connectivity, specifically, the multiplicative overhead effects and the joint effect of the autorate policy and physical layer capture.

VIII. CONCLUSION

In this paper, we performed an extensive measurement study on an operational multi-tier mesh network. We showed that the interaction of data and overhead traffic within a heterogeneous backhaul connectivity produces detrimental effects on the performance of mesh networks. We find that even low-rate overhead can produce multiplicative throughput degradation effects on data-carrying links driven by, primarily, the heterogeneity of non-forwarding links, then, by relative differences in forwarding links, and lastly, by the protocol set. Further, we find that severe throughput imbalance and aggregate throughput degradation exist between contending data flows due to a coupling of the physical layer capture effect and the misinterpretation of the channel state by the loss-based autorate mechanism.

REFERENCES

- [1] J. Camp, J. Robinson, C. Steger, and E. Knightly, "Measurement driven deployment of a two-tier mesh urban access network," in *ACM MobiSys*, Uppsala, Sweden, Jun. 2006.
- [2] A. Kamerman and L. Monteban, "WaveLAN II: A high-performance wireless LAN for the unlicensed band," *Bell Labs Technical Journal*, pp. 118–133, Summer 1997.
- [3] B. Sadeghi, V. Kanodia, A. Sabharwal, and E. Knightly, "Opportunistic media access for multirate ad hoc networks," in *Proceedings of ACM MobiCom '02*, Sep. 2002.
- [4] M. Heusse, F. Rousseau, G. Berger-Sabbatel, and A. Duda, "Performance anomaly of 802.11b," in *IEEE Infocom*, Apr. 2003.
- [5] K. K. Ramakrishnan and H. Yang, "The ethernet capture effect: Analysis and solution," in *IEEE LCN*, Minneapolis, Minnesota, Oct. 1994.
- [6] A. Kochut, A. Vasan, A. U. Shankar, and A. Agrawala, "Sniffing out the correct physical layer capture model in 802.11b," in *IEEE ICNP*, Berlin, Germany, Oct. 2004.
- [7] J. Shi, O. Gurewitz, J. Camp, and E. Knightly, "Measurement and modeling of the origins of starvation in congestion-controlled mesh networks," in *IEEE INFOCOM*, 2008.
- [8] J. Camp, V. Mancuso, O. Gurewitz, and E. Knightly, "Effects of heterogeneity in backhaul connectivity of mesh networks: A measurement study," Rice University, Electrical and Computer Engineering, Houston, TX, Tech. Rep. TREE0801, January 2008.
- [9] IEEE, "Wireless LAN medium access control and physical layer specification," ANSI/IEEE Standard 802.11, 2003.
- [10] S. Wong, S. Lu, H. Yang, and V. Bharghavan, "Robust rate adaptation for 802.11 wireless networks," in *ACM MobiCom*, 2006.
- [11] A. Iwata, C. Chiang, G. Pei, M. Gerla, and T. Chen, "Scalable routing strategies for ad hoc wireless networks," *Selected Areas of Communication*, vol. 17, no. 8, pp. 1369–1379, 1999.
- [12] K. Whitehouse, A. Woo, F. Jiang, J. Polastre, and D. Culler, "Exploiting the capture effect for collision detection and recovery," in *IEEE EmNetS*, May 2005.
- [13] C. Ware, J. Judge, J. Chicharo, and E. Dutkiewicz, "Unfairness and capture behaviour in 802.11 adhoc networks," in *IEEE ICC*, Jun 2000.
- [14] V. Kawadia and P. Kumar, "A cautionary perspective on cross-layer design," *IEEE Wireless Communications*, vol. 12, no. 1, pp. 3–11, 2005.
- [15] J. Bicket, S. Biswas, D. Aguayo, and R. Morris, "Architecture and evaluation of the MIT Roofnet mesh network," in *ACM MobiCom*, August 2005.
- [16] D. De Couto, D. Aguayo, J. Bicket, and R. Morris, "A high-throughput path metric for multi-hop wireless routing," in *ACM MobiCom*, September 2003.
- [17] S. Biswas and R. Morris, "Opportunistic routing in multi-hop wireless networks," *SIGCOMM Comput. Commun. Rev.*, vol. 34, no. 1, 2004.
- [18] D. Aguayo, J. Bicket, S. Biswas, G. Judd, and R. Morris, "Link-level measurements from an 802.11 mesh network," in *ACM SIGCOMM*, 2004.
- [19] G. Bianchi, A. D. Stefano, C. Giaconia, L. Scalia, G. Terrazzino, and I. Tinnirello, "Experimental assessment of the backoff behavior of commercial IEEE 802.11b network cards," in *IEEE INFOCOM*, Anchorage, AK, May 2007.
- [20] T. Henderson, D. Kotz, and I. Abyzov, "The changing usage of a mature campus-wide wireless network," in *ACM MobiCom*, Sep. 2004.