

Parallel Use of Multiple Channels in Multi-hop 802.11 Wireless Networks

Chen-Mou Cheng Pai-Hsiang Hsiao H. T. Kung Dario Vlah
 {doug, shawn, htk, dario}@eecs.harvard.edu
 Division of Engineering and Applied Sciences
 Harvard University
 Cambridge, MA 02138

ABSTRACT

We consider parallel use of multiple channels in a multi-radio, multi-hop 802.11 wireless network, with the goal of maximizing the total multi-hop throughput. We first quantify several fundamental forms of radio interference that cause performance degradation when the number of hops increases and that prevent total throughput from scaling up with number of radio interfaces at each node. We then evaluate three different methods of parallel channel use: Ad-Hoc, Frequency-Division Multiplexing (FDM), and Time-Division Multiplexing (TDM). We measure their performance on a linearly connected multi-hop network of dual-radio nodes. Although theoretically these three methods should have comparable performance, their actual measured performances are quite different. We find that TDM has the best performance, followed by Ad-Hoc and then FDM. The performance differences are due to these methods' capabilities of combating interference. We conclude that interference, especially adjacent channel interference, has significant effect on the achievable performance of a multi-radio, multi-hop network and hence should be carefully taken into account in the design and deployment of such a network.

I. INTRODUCTION

There has been a growing interest from the research community and industry in using low-cost Commercial Off-The-Shelf (COTS) wireless equipment, e.g., IEEE 802.11 wireless LAN ("WiFi"), for a variety of wireless

This material is based on research sponsored by Air Force Research Laboratory under agreement numbers FA8750-05-1-0035 and FA8750-06-2-0154, and by the National Science Foundation under grant number #ACI-0330244. The U.S. Government is authorized to reproduce and distribute reprints for Governmental purposes notwithstanding any copyright annotation thereon. The views and conclusions contained herein are those of the authors and should not be interpreted as necessarily representing the official policies, either expressed or implied, of Air Force Research Laboratory, the National Science Foundation, or the U.S. Government.

networking applications. In particular, there has been substantial research and deployment activity in WiFi-based wireless multi-hop networks [2], [4], [15].

However, several studies in recent years have independently shown that the performance of such networks deteriorates rapidly as the number of hops increases, due to reasons such as inefficient medium access control, radio interference, wireless link errors resulting from changing channel conditions and multipath effects, frequent route changes, and improper TCP reaction to packet loss caused by poor link quality [12], [15]. In the future, newer and longer-range technologies such as WiMAX [22] may alleviate some of the problems by using, e.g., time-division multiplexing (TDM) protocols, but it is still too early to tell how wireless multi-hop networks based on these new technologies will perform in practice.

Some of the typical approaches for enhancing the performance of multi-hop wireless networks call for use of directional antennas [14], [24], or parallel use of multiple omni-directional radios at each node [3]. These would allow a node to transmit and receive at the same time using two separate beams over the same frequency band, or, respectively, using two separate channels over two radios. Our solutions are instances of the latter approach, where each node employs two COTS radios with omni-directional dipole antennas. The approach is attractive because COTS equipment has advanced enough that it is relatively inexpensive to incorporate two radios per node, and the deployment of nodes with omni-directional antennas is relatively easy in the sense that it does not require sophisticated antenna-specific engineering. Among 802.11 COTS equipment, we are mostly interested in the newer OFDM-based 802.11a/g systems since they can deliver higher performance in terms of throughput.

The rest of this paper is organized as follows. In Section II, we identify several radio interference issues related to 802.11 multi-hop networks. Specifically, we

observe that the commercially available 802.11 chips can reliably decode packets only when the Signal to Interference plus Noise Ratio (SINR) is fairly high; for example, the OFDM-based 802.11a/g radios require at least 20dB SINR for decoding 1000-byte packets at 8% packet error rate. This decoding capability is generally worse than what people would expect. Furthermore, in environments with low path loss, there is only a slow increase of SINR with increasing interference distance. This explains why the achievable throughput in a multi-hop system decreases with the number of hops.

In Section III, we consider the case when each of the nodes in a multi-hop system is equipped with multiple radio interfaces and discuss the adjacent channel interference (ACI) issue that prevents the total throughput from scaling up with the number of radio interfaces. In Section IV, we present three methods of parallel channel use: Ad-Hoc, FDM, and TDM. In Section V, we report our performance measurement results for the three methods applied to a linearly connected multi-hop network of dual-radio nodes.

We note that our target application scenario is multi-hop networking for low-altitude unmanned aerial vehicles (UAVs), in which multi-hop forwarding is used to extend network range and provide communication beyond line of sight [5], [7]. We choose this application for two reasons. First, UAV-based networking offers many inherent application advantages, including the UAVs' capability of carrying a huge amount of on-board data relevant to the current mission and dynamically positioning themselves over regions of interest. Second, free-space inter-UAV communication suffers less from multipath complications resulting from reflections off the ground and structures on the ground such as buildings and hills. This provides us with a clean environment where inherent interference issues in a multi-hop system can be observed and analyzed more clearly. It is for the same reason of keeping the scenario as simple and clear as possible that we focus our discussion and analysis on multi-hop linear arrays.

II. INHERENT SYSTEM PERFORMANCE LIMITATIONS

Many researchers (e.g., Li et al. [15]) have examined the causes for the poor performance of multi-hop packet forwarding commonly observed in wireless ad-hoc networks. Among the many causes, the interference created by the packets of a connection when they travel along a chain of forwarding nodes imposes a stringent limit on the degree of possible spatial reuse and hence on the realizable utilization of available bandwidth. Namely, if

any two hops in a chain of k hops interfere with each other when operated simultaneously, then the realizable utilization of this chain can not be more than $\frac{1}{k}$ because there can be at most 1 active hop at a time, and each packet has to traverse at least k hops. Li et al. also verified this reasoning with both simulation and experimentation and found a close match between the two, suggesting that this $\frac{1}{k}$ -limit is probably the single most important factor contributing to the poor performance of multi-hop forwarding.

A popular model used in the ns-2 network simulator dictates that the interference distance is typically about twice that of the maximal effective communication distance; in this case, k equals to 4. However, in an outdoor experiment using ground nodes at about 20-inch elevation that we have conducted in a yard near our office building in Cambridge, MA, we observe that the interference can cover as much as four times the maximal effective communication distance for our commercially available 802.11a equipment. This makes k at least 6, rendering efficient multi-hop forwarding almost impractical for large numbers of hops. We note that our finding agrees with that of Padhye et al., who also reported the inaccuracy of the ns-2 model when experimenting with commercially available hardware [18].

We look closely into the reasons why the interference can travel so far by setting up a laboratory experiment. The experiment is designed to determine the minimum SINR required for our testbed hardware to successfully decode packets. We set up two testbed nodes, each equipped with a Wistron Neweb CM9 mini-PCI network adapter (based on the Atheros AR5213A 802.11a/b/g chipset) and use one of them as the transmitter (Tx) and the other as the receiver (Rx). We create a noise source using an Agilent 33250A 80MHz function/arbitrary waveform generator and a Maxim MAX2820 2.4GHz zero-IF transceiver: Maxim MAX2820 up-converts the baseband noise generated by Agilent 33250A to the appropriate 802.11g passband. We then connect Tx, Rx, and our noise source with a signal combiner that combines the Tx signal and the generated noise into Rx, as depicted in Figure 1.

We modify the `madwifi` driver of the Atheros chipset to let Tx broadcast 1000-byte packets continuously at the link speed of 6Mbps on channel 6 (2.427–2.447GHz). We then measure the number of packets received at Rx. We adjust the output gain of the Maxim MAX2820 transceiver until we have a packet error rate of approximately 8%. At this point, we unplug the cable going into Rx and plug it into a spectrum analyzer. The measured

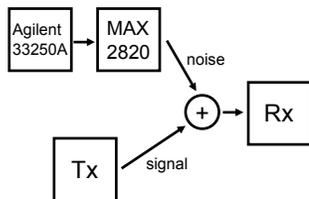


Fig. 1. The laboratory experiment setup for measuring the required SINR in order to decode 6Mbps 802.11g packets with a packet error rate of about 8%

Tx channel power is -42.5dBm , whereas the channel power of the noise source is -62.0dBm ; in other words, our testbed hardware requires about 19.5dB SINR in order to achieve a packet error rate of about 8%.

Contrary to the general belief (e.g., [20]) and to the specification of the Wistron Neweb CM9 network adapter, the SINR required for 802.11g to achieve a packet error rate of 8% at 6Mbps is not merely a few dBs—but as large as 20dB ¹. This means that the interference distance can be as large as 10 times of carrier sensing range ($k=12$), assuming a free-space propagation model, or 3 times ($k=5$) assuming a lossy-terrain propagation model. Clearly, the multi-hop throughput will be poor if k is so large.

Because of the infeasibility of sustaining reasonable throughput with a long chain of forwarding nodes, we argue that one should focus on shorter chains when aiming to achieve a high throughput in a wireless ad-hoc network. If one needs to transfer a large amount of data over a long range, one should probably traverse a small number of hops, terminate the connections at some intermediate buffering nodes, and restart the connections when appropriate. This approach is similar in spirit to the Delay Tolerant Networks framework [6]. Following this reasoning, we shall perform our experiments in a linear topology consisting of no more than four or five nodes when we investigate various techniques of exploiting parallel channels with commercially available 802.11 radios.

III. ISSUES WITH MULTI-RADIO METHODS

We examine a number of issues we have encountered in experimenting with multi-radio, multi-hop networks. We loosely divide them into issues unique to multi-radio nodes and those that apply to any multi-hop system.

¹We note that this 20dB decoding margin still conforms to the IEEE 802.11g standard [11], c.f. 17.3.10.1, Table 91—assuming a -102dBm noise floor at the antenna connector in room temperature.

A. Issues Specific to Multi-Radio Nodes

Rx-Tx Adjacent Channel Interference. The *Rx-Tx ACI* occurs when a multi-radio node transmits on one channel and, due to imperfect transmit filters in radio hardware, outputs part of the RF power into a second (adjacent) channel that another radio happens to be receiving on at the same time.

Other researchers have already encountered the ACI problem and found that combating the interference generally requires isolating the affected radios [1], [19]. For example, Robinson et al. achieve an improvement after increasing the antenna separation and shielding the radio cards [19].

In a separate paper, we reported measurements of the ACI interference using a spectrum analyzer [8]. We include some of those measurement results in Table I: the values represent relative interference power in 802.11a channels adjacent to channel 52. From this data we can see, for example, that a transmitter on channel 52 would introduce interference power to the adjacent channel 48 about 27dB lower than its main signal. Unfortunately, there is very little path loss to a co-located radio; even reduced by 27dB , the interference power is still significant, especially when competing with a faint sender.

TABLE I

Channel 52's relative interference power in the adjacent channels of 802.11a

Ch.	36	40	44	48	52	56	60	64
dB	-59	-59	-52	-27	0	-26	-53	-57

Tx-Tx Adjacent Channel Interference. The *Tx-Tx ACI* occurs when the out-of-band power from one transmitting radio is mistakenly recognized as an active carrier at another radio attempting to transmit on a different channel, causing the latter radio to back off. This prevents a node from transmitting simultaneously on both radios, even though they supposedly operate on orthogonal channels. In contrast to Rx-Tx ACI, the severity of the Tx-Tx ACI only depends on the isolation between radios on the host node; it is independent of the distance to the destination node. Lastly, the effect of the Tx-Tx ACI could possibly be avoided altogether by disabling the carrier sensing mechanism.

We measured the effect of Tx-Tx ACI by emitting a broadcast stream from two radios on a single node. The radios were set to adjacent 802.11a channels (20MHz between center frequencies) and operated at 24Mbps. We varied the offered load from 0–21Mbps, for several

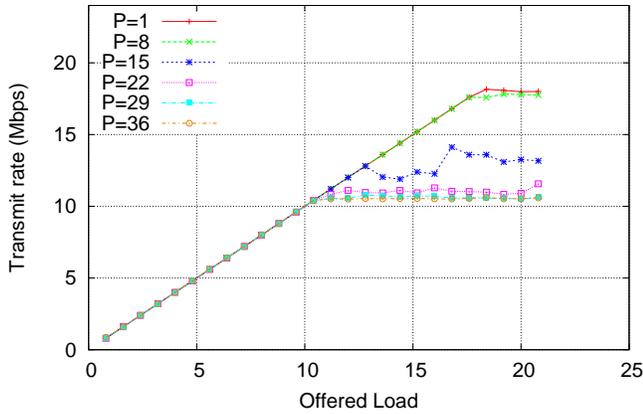


Fig. 2. Effect of Tx-Tx ACI on carrier sensing behavior (24-Mbps modulation). Transmit power levels (P) are given in Atheros driver units

different power settings. We report the observed output rates of one of the radios in Figure 2. We can see that with any power setting at or above 22, the achieved output rate reaches a peak about half of that achieved with power of 8 or lower. This is because at the higher powers, the carrier sensing mechanism gets triggered by the other, out-of-band transmission. For power settings between 8 and 22, it seems that the carrier sensing behavior is erratic; we ascribe it to the way in which the Atheros radios select the carrier sensing thresholds.

B. Issues in Single- or Multi-radio Multi-hop Systems

Co-Channel Interference from Hidden Terminals beyond Carrier Sensing Range. As we have detailed in Section II, it is possible for radios on the same channel to interfere each other even well outside their carrier-sensing range. Such a phenomenon is relatively well known in the literature; for example, Padhye et al. [18] examined this type of co-channel interference in an ad-hoc network by having all combinations of two nodes transmit at full rate and measuring the received rates everywhere else. This type of interference is the source of the well-known hidden-terminal problem [21]. Although the virtual carrier sensing mechanism (RTS/CTS) is supposed to alleviate it, we found it not quite effective in our preliminary tests; therefore, we decided to disable it in the rest of the experiments. Other researchers have also reported the ineffectiveness of RTS/CTS [15], [23].

Load Collapse on Marginal Links. We observed that when two nodes have a marginal link, increasing the transmitter’s offered load resulted in a higher packet loss rate at the receiver. We measured this effect by

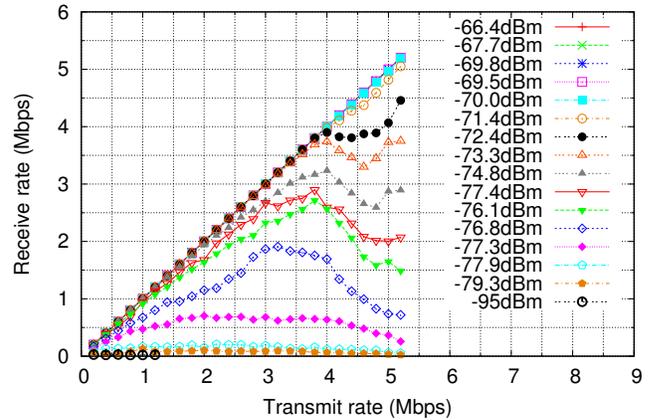


Fig. 3. UDP broadcast performance of a marginal link for a range of signal strengths and offered loads

having one node broadcast UDP packets at loads 0–6Mbps and 16 different power levels. We placed a receiver far enough to achieve a marginal link rate and recorded the incoming data rates as well as signal strengths. The results are shown in Figure 3; the 16 different power levels are identified by the average RSSI at the receiver. To rule out any problems with the sender, we placed a second receiver nearby and confirmed that it always received the full data rate. Note the sharp degradation around 4Mbps at some of the higher power configurations. We hypothesize that this effect might be a result of the specific way the Atheros radios set their receive sensitivity threshold. However, we have not been able to confirm this with Atheros engineers at this time.

These measurements show that the packet loss rate does not only depend on received signal strength, but also other transient factors; therefore, the assumptions of independent or identically distributed packet errors often used in wireless channel models may not be valid. For example, suppose we assumed that bit errors were independent. We might observe the 0.4 Mbps received for 1 Mbps sent at the -77.3 dBm power level, and conclude that the loss rate is 60%. This would lead us to predict that we should receive 2 Mbps when we send 5 Mbps across the same wireless channel. However, as we can see, the load response curve is distinctly non-linear, so the independence assumption is invalid.

To make sure that this phenomenon did not have a significant effect on our experiments, we first verified that each hop was able individually to achieve the full UDP throughput; thus, it operated above the marginal signal region which exhibits the problem.

IV. THREE METHODS OF PARALLEL CHANNEL USE

We consider the following three methods of parallel channel use: (1) Ad-Hoc, (2) Frequency-Division Multiplexing (FDM), and (3) Time-Division Multiplexing (TDM). In general, any good parallel-channel forwarding scheme should have the following properties:

- 1) **Infrequent collisions and low interference:** Reducing collisions and interference can be done at protocol level via the use of CSMA/CA and RTS/CTS. Alternatively, we can employ TDM by explicitly scheduling nodes' transmission so that no two nodes in the proximity transmit at the same time. Lastly, with FDM, we can allocate non-overlapping channels for each communication link. It is the main theme of this paper to determine experimentally how well these three methods perform.
- 2) **Aggressive spatial reuse:** Although the standard CSMA/CA and RTS/CTS mechanisms in 802.11 help reduce collisions and interference, they may lead to lower throughput by being overly conservative. Successful packet forwarding may still happen in the presence of a moderate amount of interference if the SINR is high enough (so called "capture effect"). To maximize throughput, a radio channel should be reused when the interference does not cause major performance degradation. In other words, we should seek spatial reuse as long as the link performance is acceptable in the presence of interference, rather than avoid it completely. For example, suppose we have a linear network of five nodes, A, B, C, D, and E, and we want to use the TDM method. Suppose each link can achieve 4 Mbps of throughput individually. If we use four time slots to completely avoid interference between links, we will have 1 Mbps end-to-end throughput. If, say, by operating A-B and D-E links simultaneously, we can obtain 0.7 Mbps on each link, then if we use three time slots, we will be able to achieve a higher end-to-end throughput by scheduling these two links in a time slot that is $1/0.7$ times longer than the other two time slots; the resulting throughput is $4/(1/0.7+2)=1.17$ Mbps.

In this section we describe the three methods and discuss interference issues that each method is likely to experience. A summary of each method's issues is presented in Table II.

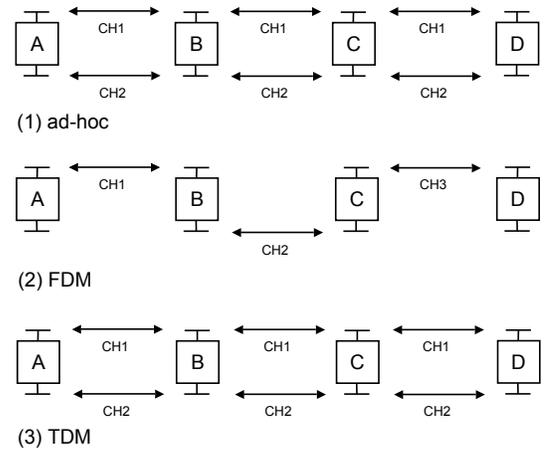


Fig. 4. **802.11 channel assignment examples for the three methods under comparison**

Let us introduce an exemplary network that consists of four nodes, as shown in Figure 4. Each of the four nodes (labeled A, B, C, D) is equipped with two radios (R_1 and R_2). The four nodes form a linear topology and are deployed so that any given node only has good links to its immediate neighbors, while its links to other non-adjacent nodes are intermittent at best. In our measurement tests, the network routes packets originating from node A to node D via B and C.

A. Ad-Hoc

Under the Ad-Hoc method, R_i on each node is assigned to 802.11 channel CH_i , for $i = 1, 2$, and multiple-access contention is resolved via the 802.11 CSMA/CA mechanism. In other words, we create two parallel ad-hoc networks, each enlisting one of the two radio interfaces from each node.

The two radios on each node can operate simultaneously; however, they share the channels with neighboring nodes using the CSMA/CA mechanism. Parallel use of the two radio interfaces this way is subject to the hidden-terminal problem and ACI (both Rx-Tx and Tx-Tx).

Hidden-terminal problem can occur when node C is outside of A's carrier-sensing range. More specifically, node A's packet transmission to B will not be received when C is also transmitting because C interferes with B's reception. We note that the RTS/CTS mechanism can only partially alleviate this problem; for example, it does not help when A interferes with D's reception from C because A and C are unlikely to be able to exchange RTS/CTS packets.

The Rx-Tx ACI occurs when one interface is receiving a packet while the other interface transmits—be it a

data packet, or a control packet like link-layer ACK. A possible way to reduce the Rx-Tx ACI while still maintaining the simplicity of ad-hoc channel access could be to “bond” the multiple radios by synchronizing all packet transmissions. In this way, any given node would likely either only transmit or only receive. We have not pursued this mechanism further, as it seemed that carrier sensing could disrupt the synchronized transmissions, and we did not have a means to disable it.

We note the severeness of both the hidden-terminal problem and Rx-Tx ACI increases with the load being transmitted at individual nodes. For the simple linear topology and single flow case, the interference may not be at its worst because the load is light at later hops. This is because the load at later hops depends on packets being delivered through the earlier hops. Thus, as some of the packets are lost due to interference, the load on later hops becomes lighter. We provide a more detailed discussion and a simple analysis of this point in Section VI.

B. Frequency-Division Multiplexing

Under the FDM method, each communication hop is assigned a unique channel. For example, node B uses channel CH1 on radio R_1 to communicate with node A and uses channel CH2 on radio R_2 to communicate with node C. With such a channel assignment, there is no need for multiple-access contention resolution such as CSMA/CA because there is a dedicated channel for each communication link. Ideally, there would be sufficiently many available channels, or enough path loss to allow reusing channels after a few hops.

In our example, the two radio interfaces on a node operate simultaneously—one for receiving packets from the upstream neighbor while the other for transmitting packets to the downstream neighbor. Clearly, use of multiple radio with this method will suffer from Rx-Tx ACI.

C. Time-Division Multiplexing

Under the TDM method, channels are assigned in the same way as in the Ad-Hoc method. The difference is that multiple-access contention now gets resolved through explicit scheduling. More specifically, we divide time into fixed-length time slots and let nodes transmit using both radios in their designated time slots. For the four-node network example we need three time slots to avoid contention. The number of time slots can be smaller than the number of hops if there are spatial reuse opportunities.

The two radio interfaces operate simultaneously during their designated time slot. Parallel channel use under TDM experiences Tx-Tx ACI and, for unicast traffic, Rx-Tx ACI due to the presence of ACKs. Because ACKs are not essential to our study of performance, we send broadcast packets instead of unicast packets so as to disable link-layer ACKs.

For TDM to work, nodes need to be synchronized in time, so all nodes can follow the schedule correctly. In addition, there is some queuing inside the operating system software that must be accounted for in order to properly schedule the transmissions.

Timing errors can reduce efficiency of parallel channel use due to overlapping time slots: during the overlap, the affected nodes either back off due to carrier sensing, or lose the data due to interference. With longer time slots, we can tolerate higher timing error. For example, 1 ms timing error per 100-ms time slot results in no more than 2% efficiency loss (1% at the beginning and 1% at the end of a time slot may overlap with adjacent time slots). But increasing the time slot length increases the end-to-end latency and the buffering requirements on each forwarder. This limits the size of the time slot to be a few hundred milliseconds in most practical applications.

There are several timing synchronization protocols that can achieve sub-millisecond accuracy, for example, Reference-Broadcast Synchronization (RBS) [9] and Continuous Clock Synchronization [17]. In our experiments, we manually synchronize our nodes using `ntpdate` before each experiment to achieve millisecond timing accuracy. Note that `ntpdate` is an NTP client that only synchronizes local clock’s phase offset to a common server, and it does not continuously track timing error for local clock’s frequency offset. Because our experiments only last a short period of time (several minutes at most), local clock’s frequency offset has little effect on the timing accuracy.

The operating system employs packet queues to minimize overhead caused by context switching; often, these queues can buffer a large number of packets. For example, UDP sockets on Linux machines have a default queue size of 100KB, translating into more than 60 packets in our setting. Furthermore, the device driver uses a queue to amortize interrupt processing overhead. In the `madwifi` driver we use in our experiments, the default queue size is 200 packets. Given that the transmission timing control runs as a user process, packets stored in these two queues would continue to drain outside of the designated time slot even after the user process stops transmitting. To minimize the effect

TABLE II
Expected interference issues for each method

	Rx-Tx ACI	Tx-Tx ACI	Hidden Terminal
Ad-Hoc	√	√	√
FDM	√		
TDM	√ (ACK only)	√	

of queuing, we reduce these two queue sizes to 3 and 4 packets, respectively. In addition, we program the user process to stop sending data early to compensate for the transmission delay of queued packets (e.g., the application stops 14 ms before the end of the time slot if each packet takes 2 ms to transmit).

V. MEASUREMENT RESULTS

We conducted a set of experiments to measure the throughput achieved by the three multi-radio, multi-hop methods described in Section IV. We ran the experiments on a testbed network consisting of 400MHz AMD Geode single-board computers made by Thecus Inc. We equipped each of these nodes with two Wistron Neweb CM9 mini-PCI network adapters, which are based on the Atheros AR5213A 802.11a/b/g chipset. The network adapters ran the `madwifi` Linux driver². We used a mix of two types of antennas: (1) off-the-shelf, dual-band, 7dBi omni-directional antennas, and (2) 2dBi dipole antennas made in-house.

The experiments were conducted outdoors in a yard outside our office building as well as in a large athletic field across the campus. Our testbed consisted of five nodes, A, ..., E, which were arranged in a linear topology similar to that depicted in Figure 4. We set the transmit power equally on all nodes, and then placed the nodes on the ground so that each hop was as large as possible while still maintaining low-loss communication. We performed the 802.11a measurements in the yard outside our office building where we did not have any 802.11a networks and the 802.11b/g measurements in the athletic field where there were no other 802.11b/g networks.

All links ran at the lowest possible modulation rates—6Mbps for 802.11a/g and 1Mbps for 802.11b. Measured single-hop throughputs for these three modulations are 5.6Mbps, 4.4Mbps, and 0.9 Mbps, respectively, calculated from UDP payload based on 1472-byte packets.

²A newer version of the driver, called `madwifi-ng`, was available, but we didn't use it due to stability issues.

In all three multi-radio methods, a user-level source process at node A sent 1472-byte UDP packets toward a sink process at node E. In Ad-Hoc and FDM methods, we used a 100KByte kernel-level send buffer; when this buffer became full, the source process would be blocked from sending more packets. In contrast, in the TDM method, we buffered packets in a large user-level queue and used a minimal 7-packet kernel-level send buffer (3 in the socket buffer plus 4 in the hardware transmit queue) to avoid spillover outside the designated time slot. We have made sure that the buffer occupancies were reasonable in all three methods so that they would not have a significant effect on the throughput measurement result.

We divide the results into three subsections below: 1) comparing TDM, Ad-Hoc, and FDM on a three-hop chain, 2) comparing three- and four-slot TDM on a four-hop chain, and 3) comparing 802.11b and g. In all cases, the measurements are listed as three values: the measured end-to-end UDP throughput, the ideal rate achievable assuming orthogonal channels and zero packet-loss probability, and a ratio of measured to ideal rate for ease of comparison.

We calculate the ideal rates from single-hop capacities as follows. Let C be the single-hop capacity. For Ad-Hoc and TDM, we assume that the k hops (where k is either 3 or 4) are all in each other's interference range, and hence share the channel. For two-radio nodes, there are two independent k -hop forwarding chains where each hop gets C/k Mbps, and so the maximum achievable throughput is $2 \cdot C/k$ Mbps. With FDM on three hops, all hops can operate simultaneously at C Mbps, making that the end-to-end throughput. With FDM on four hops, we assume the first and last hop evenly share the same channel, leading to a bottleneck rate and ideal throughput of $C/2$.

A. Comparison of TDM, Ad-Hoc, and FDM over three hops

802.11a:

	Achieved	Ideal	Percentage
FDM	2.83 Mbps	5.6 Mbps	51%
Ad-Hoc	2.50 Mbps	3.7 Mbps	68%
TDM	3.30 Mbps	3.7 Mbps	89%

802.11g:

	Achieved	Ideal	Percentage
FDM	1.87 Mbps	4.4 Mbps	43%
Ad-Hoc	1.94 Mbps	2.9 Mbps	67%
TDM	2.78 Mbps	2.9 Mbps	96%

B. Three- vs. four-slot TDM over four hops

802.11g:

	Achieved	Ideal	Percentage
TDM-3	1.66 Mbps	2.9 Mbps	57%
TDM-4	2.00 Mbps	2.2 Mbps	91%

C. 802.11g vs. 802.11b over four hops

As stated above, since 802.11b and g have only 3 non-overlapping channels, at least one channel must be reused. In this case, we put the first and last hop on the same channel, leading to an ideal throughput of $C/2$. We note that Ad-Hoc performs quite well when 802.11b is used, possibly because 802.11b has a lower decoding margin, making it more resilient to interference.

802.11g:

	Achieved	Ideal	Percentage
FDM	1.40 Mbps	2.8 Mbps	50%
Ad-Hoc	1.49 Mbps	2.2 Mbps	68%
TDM	2.00 Mbps	2.2 Mbps	91%

802.11b:

	Achieved	Ideal	Percentage
FDM	0.34 Mbps	0.45 Mbps	76%
Ad-Hoc	0.37 Mbps	0.45 Mbps	82%
TDM	0.36 Mbps	0.45 Mbps	80%

VI. DISCUSSION

In this section, we report the measurement result that demonstrates the impact of interference from hidden terminals beyond carrier sensing range, as described in Section III. Specifically, we compare the performance of a 3-slot TDM scheme for a 3-hop network (3HOP) and a 4-hop network (4HOP). These two cases are interesting because they are otherwise the same configuration, except that in 4HOP, the first hop and the last hop share the same time slots and hence will be transmitting simultaneously. Our initial 802.11g measurement result shows that the transmitters of these two hops can not sense each other's transmission; in other words, they both transmit at full rate as if there were no other transmitters in their neighborhood. Ideally, if there were no interference beyond carrier sensing range, then we would observe close, if not identical, performance results for 3HOP and 4HOP; however, our measurement result in Subsection A and B of the previous section shows that 4HOP's performance is only 60% of 3HOP's (1.66Mbps vs 2.78Mbps).

There is a simple calculation that we can do to quantitatively determine the impact of such interference.

We first experimentally determine the effect of such interference by operating the two links simultaneously with full load. In this example, the interference is quite asymmetric: the first hop suffers much more than the last hop, most likely because the transmitter of the last hop is closer to the receiver of the first hop than the other way around. Specifically, when both links are busy, we get 30% throughput on the first hop and 96% on the last hop. Now in the TDM scheme, all packets originate from the first hop; if the transmission of packets on the last hop kills some of the packets on the first hop, then subsequently the last hop will get less packets, resulting in underutilization of the last hop and hence less interference for the first hop. In our experiment setup, all intermediate hops have reasonably good performance, so we are going to simplify the analysis by assuming that they are lossless; the same simplification can be applied to the last hop since its performance is fairly close to 100% even in the presence of interference from the first hop. Now call the probability of the transmission on the last hop killing the transmission on the first hop α (in this case, $\alpha = 0.7$). In equilibrium, suppose the percentage of packets getting through the first hop is x ; this is going to be the utilization of the last hop since we assume the intermediate hops are lossless. These packets are going to kill αx of the packets on the first hop, making the throughput of the first hop $1 - \alpha x$, which must be equal to x by conservation of traffic flow. This means that $x = (1 + \alpha)^{-1}$, or 0.59 in our experiment, which is fairly close to the measurement result of 60%.

VII. CONCLUSION

In this paper, we considered the problem of efficient parallel use of 802.11 channels in multi-radio, multi-hop wireless networks. We experimentally evaluated three methods of parallel channel use, based on common MAC schemes (Ad-Hoc, FDM, TDM), on a linearly connected dual-radio network. The two main contributions of this paper are: 1) the finding that ACI is a significant obstacle for good performance in multi-radio networks; and 2) the proposed solution for mitigating the ACI effect using TDM channel access, supported by experimental evidence.

Among the three methods, TDM performs the best in all the experiments except for the four-hop case of 802.11b, in which Ad-Hoc performs slightly better than TDM. This is because TDM avoids most types of interference by design. By dividing channel uses into separate time slots, TDM schedules better than Ad-Hoc, but the two perform pretty closely when only a single

channel is used. Ad-Hoc, despite being the simplest method among the three, performs better than FDM because it manages to avoid some of the Rx-Tx ACI in FDM.

For nodes with more than two radios, we believe that among the three methods, only TDM's performance can scale up as the number of radios increases. The other two methods do not explicitly address the ACI issue, and so with more radios, the effect of ACI is likely to get worse.

In summary, the results of this paper show the importance of considering ACI in protocol design for multi-radio networks. Otherwise, as we have seen in the case of Ad-Hoc and FDM, real-world systems would not perform as well as expected. Thus, analysis taking into account ACI is required to design and understand the performance of multi-radio protocols properly.

REFERENCES

- [1] Adya, A., Bahl, P., Padhye, J., Wolman, A. and Zhou, L., "A Multi-Radio Unification Protocol for IEEE 802.11 Wireless Networks," First International Conference on Broadband Networks (BROADNETS'04), 2004
- [2] Aguayo, D., Bicket, J., Biswas, S., Judd, G., and Morris, R., "Link-level Measurements from an 802.11b Mesh Network," *SIGCOMM 2004*, August 2004
- [3] Alicherry, M., Bhatia, R., and Li, L., "Joint Channel Assignment and Routing for Throughput Optimization in Multi-radio Wireless Mesh Networks," *ACM MobiCom 2005*, August 2005
- [4] Bicket, J., Aguayo, D., Biswas, S., and Morris, R., "Architecture and Evaluation of an Unplanned 802.11b Mesh Network," *ACM MobiCom 2005*, August 2005
- [5] Brown, T. X., Argrow, B., Dixon, C., Doshi, S., "Ad Hoc UAV Ground Network (AUGNet)," *AIAA 3rd "Unmanned Unlimited" Technical Conference*, Chicago, IL, 2004.
- [6] Cerf, V., Burleigh, S., Hooke, A., Torgerson, L., Durst, R., Scott, K., Fall, K., and Weiss, H., "Delay-Tolerant Network Architecture," Internet Draft, March, 2006.
- [7] Cheng, C.-M., Hsiao, P.-H., Kung, H. T., and Vlah, D., "Performance Measurement of 802.11a Wireless Links from UAV to Ground Nodes with Various Antenna Orientations," *ICCCN 2006*, Arlington, VA, October, 2006.
- [8] Cheng, C.-M., Hsiao, P.-H., Kung, H. T., and Vlah, D., "Adjacent Channel Interference in Dual-radio 802.11 Nodes and Its Impact on Multi-hop Networking," *IEEE GLOBECOM 2006*, San Francisco, CA, November 2006.
- [9] Elson, J., Girod, L., and Estrin, D., "Fine-grained network time synchronization using reference broadcasts," *SIGOPS Operating System Review*, Volume 36, Issue SI, 2002
- [10] Hyperlink Technologies, Inc., "2.4 GHz 802.11b and 802.11g Compatible 4-Pole Ultra High Q WiFi Bandpass Filters," <http://www.hyperlinktech.com/web/bpf24-4xx.php>, 2005
- [11] IEEE 802.11 Working Group, "Wireless LAN Medium Access Control (MAC) and Physical Layer (PHY) specifications," *IEEE 802.11 standard, including 802.11a and 802.11b extensions*, September 1999
- [12] Jain, K., Padhye, J., Padmanabhan, V. N., and Qiu, L., "Impact of Interference on Multi-hop Wireless Network Performance," *ACM MobiCom 2003*, September 2003
- [13] Judd, G., and Steenkiste, P., "A simple mechanism for capturing and replaying wireless channels," *E-WIND '05: Proceeding of the 2005 ACM SIGCOMM workshop on Experimental approaches to wireless network design and analysis*, 2005
- [14] Ko, Y. B., Shankarkumar, V., and Vaidya, N., "Medium Access Control Protocols using Directional Antennas in Ad Hoc Networks," *IEEE INFOCOM*, March 2000
- [15] Li, J., Blake, C., De Couto, D. S. J., Lee, H. I., and Morris, R., "Capacity of Ad Hoc Wireless Networks," *ACM MobiCom*, July 2001
- [16] Mikrotik Ltd, "Routerboard 500," <http://www.routerboard.com/rb500.html>, 2006.
- [17] Mock, M., Frings, R., Nett, E., and Trikaliotis, S., "Continuous Clock Synchronization in Wireless Real-Time Applications," *Symposium on Reliability in Distributed Software*, pp. 125-132, 2000
- [18] Padhye, J., Agarwal, S., Padmanabhan, V., Qiu, L., Rao, A., and Zill, B., "Estimation of Link Interference in Static Multi-hop Wireless Networks," *Internet Measurement Conference (IMC)*, Berkeley, CA, October 2005
- [19] Robinson, J., Papagiannaki, K., Diot, C., Guo, X. and Krishnamurthy, L., "Experimenting with a Multi-Radio Mesh Networking Testbed," *Ist workshop on Wireless Network Measurements (WiNMee 2005)*, Trento, Italy, April 2005
- [20] Thomson, J., Baas, B., Cooper, E. M., Gilbert, J. M., Hsieh, G., Husted, P., Lokanathan, A., Kuskin, J. S., McCracken, D., McFarland, B., Meng, T. H., Nakahira, D., Ng, S., Ratterhalli, M., Smith, J. L., Subramanian, R., Thon, L., Wang, Y.-H., Yu, R., and Zhang, X., "An Integrated 802.11a Baseband and MAC Processor," *Proc. ISSCC 2002*, San Francisco, CA, February 2002
- [21] Tobagi, F. A., and Kleinrock, L., "Packet Switching in Radio Channels: Part II - the Hidden Terminal Problem in Carrier Sense Multiple-Access Modes and the Busy-Tone Solution," *IEEE Transactions on Communications*, 23(12):1417-1433, 1975
- [22] The WiMAX Forum, <http://www.wimaxforum.org>
- [23] Xu, K., Gerla, M., and Bae, S., "How effective is the IEEE 802.11 RTS/CTS Handshake in Ad Hoc Network," *Proc. IEEE GlobeCom'02*, Taipei, Taiwan, Nov. 2002
- [24] Yi, S., Pei, Y., and Kalyanaraman, S., "On the Capacity Improvement of Ad Hoc Wireless Networks Using Directional Antennas," *ACM MobiHoc 2003*, June 2003