Listen (on the Frequency Domain) Before You Talk

Souvik Sen
Duke University
Durham, NC, USA
ssen@cs.duke.edu

Romit Roy Choudhury
Duke University
Durham, NC, USA
romit@ee.duke.edu

Srihari Nelakuditi
University of South Carolina
Columbia, SC, USA
srihari@cse.sc.edu

ABSTRACT

Conventional WiFi networks perform channel contention in time domain. This is known to be wasteful because the channel is forced to remain idle, while all contending nodes are backing off for multiple time slots. This paper proposes to break away from convention and recreate the backing off operation in the frequency domain. Our basic idea is to pretend that OFDM subcarriers are integer numbers, and thereby, view today's random backoff process as equivalent to transmitting on a randomly chosen subcarrier. By employing a second antenna to listen to all the subcarriers, each node can determine whether its chosen integer (or subcarrier) is the smallest among all others. In fact, each node can even determine the rank of its chosen integer, enabling the feasibility of a TDMA-like schedule from every round of contention. We develop these ideas into a *Time to Frequency* (T2F) protocol and prototype it on a small testbed of 8 USRPs. Experiments confirm its feasibility, along with promising throughput gains of more than 35% at high bit rates. A fuller design and thorough evaluation of T2F is a topic of ongoing work.

Categories and Subject Descriptors

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

General Terms

Design, Experimentation, Performance

1. INTRODUCTION

Access control strategies are designed to arbitrate how multiple entities access a shared resource. Several distributed protocols embrace randomization to achieve arbitration. In WiFi networks, for example, each participating node picks a random number from a specified range and begins counting down. The node that finishes first, say N_1, wins channel contention and begins transmission. The other nodes *freeze* their countdown temporarily, and revive it only after N_1's transmission is complete. Since every node counts down at the same rate, this scheme produces an implicit ordering among nodes. Put differently, the node that picks the smallest random number transmits first, the one that picks the second-smallest number transmits second, and so on. The overall operation is often termed as "backoff".

While backoff arbitrates channel contention, it incurs a performance cost. Specifically, when multiple nodes are simultaneously backing off, the channel must remain idle, naturally leading to underutilization. Moreover, network congestion prompts exponential increase in the backoff range, introducing the possibility of greater channel wastage. Authors in [1] show more than 30% reduction in throughput due to backing off; [2] shows the severity at higher data rates. This paper attempts to address this problem by migrating the backoff operation to the frequency domain.

Our main idea is simple. When a node N_1 has a packet to transmit, it picks a random value, r_1, from a specified range $[0, F]$. Once the channel becomes idle, N_1 begins the backoff operation. However, instead of counting down from r_1 to 0, N_1 transmits a symbol on the r_1^{th} subcarrier[1]. We assume that each node has two antennas; thus, while one antenna transmits, the other antenna listens to determine which of the subcarriers are active. Assuming N_2 is also contending for the channel, and say has transmitted on the r_2^{th} subcarrier, N_1 observes activity on both subcarriers r_1 and r_2. Assuming $r_1 < r_2$, N_1 immediately infers that it has won channel contention, and begins transmission[2]. N_2 learns that it has lost, and defers its own transmission until N_1 has finished. We call this approach *T2F*.

[1]Subcarriers are narrowband OFDM channels used by 802.11.
[2]Section 3.1 describes how T2F uses another round of subcarrier based contention to avoid a collision (two nodes pick the same smallest random number). However, for ease of illustration, here we present T2F as if it resolves contention with one round.

Hotnets '10, October 20–21, 2010, Monterey, CA, USA.

The advantages of T2F are two fold. First, one round of frequency domain backoff lasts for one OFDM symbol, far less than the average backoff in protocols like 802.11. Second, T2F creates a logical ordering among contending nodes, and each node learns its own rank in this order. This ranking among nodes creates the possibility of back-to-back transmissions – a short TDMA schedule – eliminating the need for per-packet backoff. Since 802.11 currently backs off before every packet, T2F helps improve the channel usage and network throughput.

Of course, extracting these gains entail a number of research challenges: (1) Active subcarriers need to be detected accurately, in face of loose time synchronization among transmitters and energy leakage between narrowband subcarriers. (2) Collision among nodes – which happens when multiple nodes choose the same subcarrier – needs to be mitigated successfully. (3) Once nodes determine the transmission order, T2F needs to enforce this ordering, even in the presence of interference, newly joining nodes, transmission failures, etc. We address these challenges, consolidate them into a protocol, and implement a proof-of-concept on the USRP/GNURadio platform. Experimental results show 95% accuracy in subcarrier detection, 0.03 probability of collision, and throughput gains of more than 35% over 802.11 in a network of 4 links. In summary, the contributions in this paper are as follows:

- *We identify an opportunity to migrate protocol operations from the time to the frequency domain.* Although we instantiate our ideas through a WiFi based MAC, they may be generalized to other access strategies.
- *We design an OFDM based system where random backoff is realized by selectively transmitting on a subcarrier.* A logical order among senders is enforced in a decentralized manner, for improved channel usage.
- *We address the challenges behind such a scheme, and prototype it on the USRP/GNURadio platform.* Stable behavior, along with appreciable performance gains, give us confidence to build a large, full-scale system.

2. 802.11 AND OFDM

This section highlights the limitations of 802.11's backoff (in time domain), and presents a simple abstraction of OFDM (to better explain the shift to frequency domain).

802.11 Channel Access: WiFi prescribes each transmitter to backoff for r time slots, where r is randomly chosen from the range $[0,31]$. Each time slot corresponds to $9\mu s$. The node counts down only if the channel is idle – if the node senses a signal on the channel, the countdown is frozen, and revived only after the channel is idle again. Whichever node completes the countdown first, say N_1, begins transmission. When this transmission is complete, the other nodes continue with their remaining countdown. *Observe that 802.11 implicitly forms a queue among contending nodes, each node's position in the queue determined by the random number it chooses.* Prior analyses have shown that this simple scheme guarantees stability and (long-term) fairness.

We make three observations that are not necessarily new. (1) Fundamentally, backing off is not a temporal domain operation. Its implementation is in the temporal domain, forcing the channel to be idle per-packet. (2) The duration of each backoff slot is fixed, implying that the channel wastage grows at higher data rates. This is because a packet's airtime is shorter at higher rates, and hence, the fraction of channel time occupied by idle slots is larger. (3) Although channel utilization increases with few nodes (i.e., idle slots are amortized), too many nodes cause collisions. A collision forces nodes to exponentially increase their backoff, pushing the system back to underutilization (see extensive analysis and measurements in [1,3]).

Orthogonal Frequency Division Multiplexing (OFDM) can be abstracted as a PHY scheme that divides the wireless spectrum into multiple narrow band channels, called subcarriers. The subcarriers carry modulated data streams in parallel, but at a lower rate per-subcarrier. The benefit of OFDM emerges from its ability to cope with channel adversities, including narrowband interference and frequency-selective fading due to multipath. The 802.11 implementation of OFDM has 52 subcarriers, of which 48 are used for data transmission, and 4 for equalization. A transmitter stripes bits across all subcarriers, however, it is possible to transmit/receive only on a subset of them.

As we will see later, a T2F node picks a random number, say 13, and transmits a signal *only* on the 13^{th} subcarrier. The node's second antenna detects a strong signal on the 13^{th} subcarrier, as well as on other subcarriers used by other contending nodes. Practical hardware constraints raise difficulties in discriminating between adjacent subcarriers. When the second antenna receives a strong signal on the 13^{th} subcarrier, "leakage" into adjacent subcarriers may mislead the receiver into detecting subcarriers 12 and 14 also as active. Higher sized FFTs (such as 256 point) are useful to mitigate such effects – the spikes on subcarriers can be better isolated. Widening the subcarrier width helps as well. Migrating to the frequency domain brings these problems, and T2F needs to cope with them.

3. ARCHITECTURE AND DESIGN

T2F's main modules are: subcarrier based backoff and scheduled transmission. For ease of explanation, we consider contention among multiple APs (or any wireless nodes) located within the same collision domain. Later, we discuss how T2F extends to multiple collision domains.

3.1 Subcarrier based Backoff

When an ongoing transmission ends, all APs wait for a DIFS (Distributed Coordinator Function Inter Frame Space) duration, choose a random number in the range [1, 52], and transmit a signal on the corresponding subcarrier. Figure 1 shows an example where AP1 chooses 12 and AP2 chooses 36. These subcarrier transmissions are concurrent because the APs get implicitly synchronized when the channel becomes idle [4]. The listening antenna on each AP receives the combined signals from all the APs, as well as its own signal, called the *self-signal*. The listening antenna can then detect all the active subcarriers (i.e., subcarriers on which all the APs have transmitted, including itself). A simple way of determining the winner would be to compare one's own backoff value (subcarrier number) with others. The AP with the smallest number could proceed with data transmission. In the example of Figure 1, AP1 gets to transmit next. The operation consumes one OFDM symbol time (= 3.2μs).

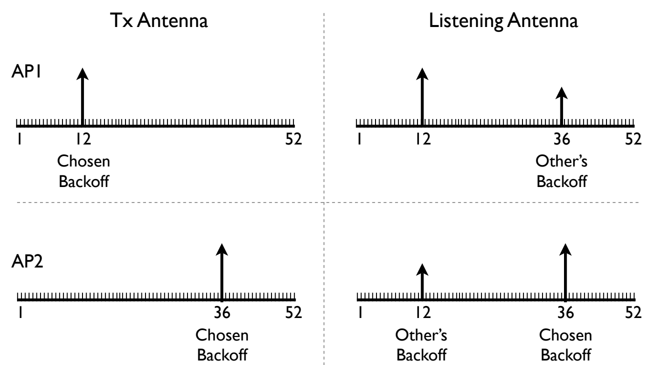

Figure 1: Backoff in the frequency domain. AP1 and AP2 randomly pick 12 and 36 respectively and transmit on the corresponding subcarriers. The listener antennas of both APs can detect that subcarriers 12 and 36 are active. Then, AP2 defers transmission in favor of AP1.

Of course, if two APs pick the same subcarrier number, they would collide. T2F can avoid this by introducing another round of subcarrier based contention. A node that believes is a winner in the first round, retransmits on another randomly chosen subcarrier immediately after. This process is illustrated in Figure 2. Both AP1 and AP2 pick 12 in the first round, but AP1 wins the contention by picking 13 (against AP2's 29) in the second round. If multiple nodes chose the minimum number in the first round, the probability of them coinciding again is small. T2F can reduce the collision probability to an arbitrarily small value, at the expense of more rounds (OFDM symbol durations).

3.2 Scheduled Transmissions

802.11's backoff mechanism does not allow nodes to estimate its rank in the order of transmissions. This is because the backoff value is not shared among nodes – it is a distributed and implicit form of ordering. T2F, however, enables each node to learn its rank in the sequence of pending transmissions. In fact, each node also knows the exact backoff values chosen by other nodes (although the mapping between node and backoff value is not known). T2F

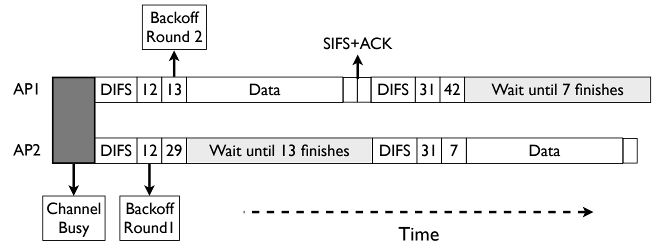

Figure 2: Two rounds of subcarrier based backoff. In the first round, both AP1 and AP2 draw the same random number 12. But in the next round, AP1 wins the contention by picking 13 against 29 by AP2. This process repeats after AP1 finishes its transmission.

exploits this knowledge to enable back-to-back, TDMA style transmissions. A node ranked n transmits immediately after the $(n-1)^{th}$ ranked node finishes transmission. The protocol structure is as follows.

As a first step, multiple nodes are promoted to the second round of contention. If N APs contend, the *top-K* ranked APs advance to the second round ($K \leq N$). The choice of K brings out a tradeoff between collision probability and TDMA schedule length. Higher values of K will result in longer TDMA length (better throughput), but at the risk of collisions. While we will revisit this tradeoff in the next subsection, we observe here that a given choice of K can be implemented. This is because after a round of contention, each node knows its rank, and can independently decide whether it should advance to the next round. For example, assume that 4 APs contend in the first round as in Figure 3. Suppose the nodes ranked in the top two are allowed to enter the second round. Both AP3 and AP4 can figure that they are not top two nodes. So only AP1 and AP2 enter the second round of contention.

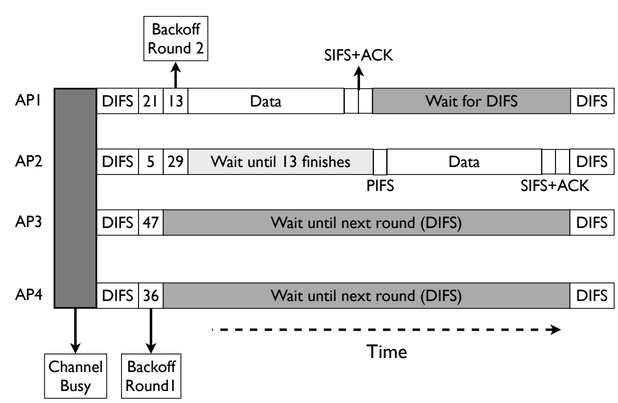

Figure 3: Backoff in the frequency domain followed by scheduled transmissions. All APs contend in the first round, but only AP1 and AP2 enter the second round. Based on backoff values in the second round, the schedule is AP1 followed by AP2. Only after the scheduled transmissions complete, AP3 and AP4 contend again.

The top-K nodes are reordered after the second round of contention – this is the order of data transmissions. To enforce this order, each AP includes its own backoff value in the PLCP (Physical Layer Convergence Protocol) header of its data packet. The other APs overhear these values and wait until the AP just ahead in sequence has finished transmission. At that time, the AP waits for a PIFS (Point Coordinator Function Inter Frame Space) duration (shorter than DIFS) and starts transmitting. In our example, AP1 includes its backoff value 13 in the PLCP of the data packet. AP2 waits until it has heard a packet with 13 as its backoff; once this packet completes, AP2 initiates transmission. Since all the APs are within a single collision domain, and since PLCP headers are detected with high reliability, the APs transmit back-to-back in the order AP1, AP2 (a simple form of TDMA).

3.3 Points of Discussion

(1) How does T2F extend to multiple collision domains? During the TDMA-type schedule, a node waits for the PLCP header from the node just ahead of it in the schedule. In addition, it senses the channel and initiates transmission only if the channel is idle for a PIFS duration. We believe this allows multiple collision domains to coexist. Suppose AP1, AP2, and AP3 belong to one collision domain (see Figure below), and AP2 and AP4 belong to a different collision domain (i.e., AP4 does not carrier sense AP1 or AP3). Also, say the transmission order has been decided as {AP1, AP2, AP3}. When AP1 is in progress, it is possible that AP4 initiates transmission. Since AP2 senses AP4, it defers its own transmission; now, since AP3's transmission is predicated upon AP2's, AP3 also defers. The relative ordering among AP1, AP2, and AP3 is maintained, even with overlapping collision domains. Further analysis and evaluation is required to validate our intuition.

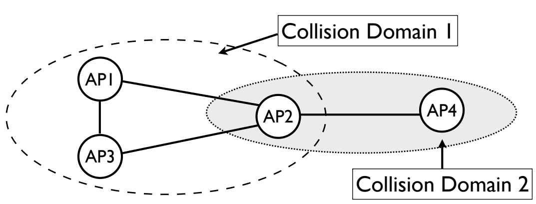

(2) How does T2F choose the value of K, the number of nodes promoted to the second round of contention? As mentioned earlier, higher K increases the length of the TDMA schedule, but at the risk of collisions. T2F handles this by computing the variation of collision probability for increasing K, and translating it to net throughput. Observe that the collision probability is not for the winning node alone; to avoid any collision, all nodes in the second round should pick unique subcarriers. We omit the analysis in the interest of space, but only mention that in practical settings with 52 subcarriers, $K = 3$ is effective. Increasing the number of subcarriers is also a control knob. However, more subcarriers mean narrower bandwidth per-subcarrier, which in turn implies susceptibility to leakage and detection error. A full treatment of this tradeoff is part of our future work.

(3) Transmissions on subcarriers are not tightly time synchronized – how does this impact T2F? While lack of synchronization is an issue, transmitting on the subcarriers for slightly longer mitigates the problems. To elaborate, nodes perform FFTs to determine the active subcarriers. Each node requires at least one FFT window within which all the active subcarrier signals are present. The staggering between these signals (i.e., their lack of synchronization) originates from two main sources. (1) When the channel changes from busy to idle, different nodes detect it at different times because of unequal propagation delay. (2) Once all nodes begin their backoff transmission, the signals arrive at any receiver with some separation, caused by propagation delay again. Assuming a maximum propagation delay of t_{pd} between any two nodes within a collision domain, the stagger between two signals should be bounded by $2t_{pd}$. Typical propagation delays in WLANs are within $1\mu s$ [5]. Since a $64pt$ FFT takes 3.2μs at 20MHz, each contention round must extend for $5.2\mu s$. Together, two rounds of T2F contention incurs $10.4\mu s$, considerably smaller compared to $150\mu s$ backoff on average in 802.11.

(4) The self-signal from the transmitting antenna to its own listening antenna is strong – how does this affect the detection of other subcarriers at the listening antenna? Figure 4(a) shows the FFT spectrum at the listening antenna when the transmit antenna transmits on subcarrier 2. The x-axis lists the subcarrier numbers from our USRP prototype. The outcome is a high impulse around subcarrier 2 and a naturally high DC component (near X=0). Like 802.11, T2F also skips subcarriers around 0 to avoid the DC effects. Thus, the self-subcarrier is clearly discernible. We also look at the capability to discern multiple subcarriers when they are adjacent – we make the inter subcarrier spacing equal to 802.11 (i.e., $0.3125 MHz$). Figure 4(b) confirms the feasibility with a 256 point FFT at the listener. Even in presence of a high self-subcarrier, transmissions on various other subcarriers are clearly discernible.

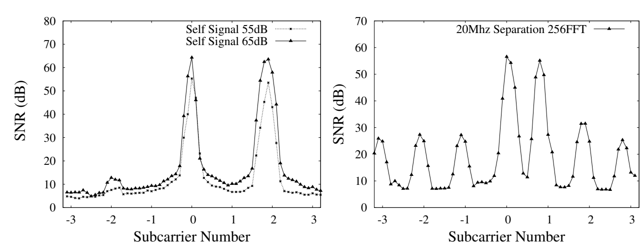

Figure 4: Active subcarrier detection: (a) Self signal on subcarrier 2; (b) 6 nodes transmitting at subcarrier distances corresponding to 20Mhz, 256pt FFT listening antenna. All active subcarriers are clearly discernable.

(5) How does the SNR of the subcarrier signals affect T2F? In other words, will weak signals from far away nodes be discernible, especially in the presence of self-signal? Figure 5(a) shows the detection accuracy of subcarriers with varying SNRs and increasing separation from the strong self-subcarrier. At $15dB$ or more, the detection accuracy is more than 95%, even when another node's chosen subcarrier is immediately adjacent to the self-subcarrier. Spectral separation from the self-subcarrier further increases the accuracy. To show that poor detection accuracy at lower SNR is not due to the strong self-signal, we performed experiments without the self-signal. Figure 5(b) shows that even then, USRP hardware is unable to discern the subcarriers reliably at $10dB$ SNR. This implies that the impact of self-signal on subcarrier detection is marginal.

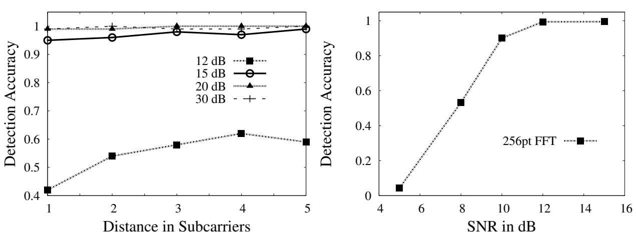

Figure 5: Detection accuracy: (a) as a distance from the self subcarrier with 256pt FFT at the listening antenna; (b) in the absence of strong self-subcarrier.

4. PROTOTYPING AND EVALUATION

Two obvious questions about T2F are whether it is feasible and if so what is the gain. To this end, we first describe our prototype, and then evaluate the feasibility of detecting subcarriers. We then conduct simulations to evaluate collision probability and T2F's throughput gain over 802.11.

USRP/GNURadio Prototype: We prototype T2F on a small testbed of 8 USRPs to evaluate the subcarrier detection accuracy. The distance between the transmitting and listening antenna is around 20 inches resulting in a strong self-signal of 55dB [3]. The transmitter transmits in 2MHz band, while the listening antenna samples the 8MHz channel, both around the same central frequency. This allows the receiver to detect a subcarrier within one OFDM symbol duration by employing an FFT four times that of the transmitter, i.e., if the transmitter uses 64 point IFFT the receiver uses 256 point FFT. Note that this requirement is specific to USRPs since the current 802.11a/g OFDM designs on USRPs need higher FFT sizes due to imprecision [7]. We believe commercial hardware has greater precision, removing the need of higher FFT sizes at the receiver.

The listener antenna detects subcarriers using a joint thresholding and peak-detection scheme. This is necessary because with practical hardware (especially USRPs), the subcarriers do not emerge as impulses, but are instead like peaks (see Figure 4(a)). T2F declares a peak whenever it is above a chosen threshold, determined as a function of the periodically sampled noise floor [3, 5]. Our experiments show that a threshold between 7 to 12dB SNR is adequate to achieve high detection accuracy.

Subcarrier detection: As a first step, T2F transmitters are made to transmit signals on randomly chosen subcarriers; the listening antenna performs the FFT and detects the subcarrier. With 802.11-like subcarrier separation (.3125MHz), 2MHz band amounts to 6 subcarriers. We have also experimented with 24 much narrower subcarriers. *Note that the choice of subcarriers is limited by the processing capability of USRP hardware not a limitation of T2F.* Figure 6(a) and (b) show the false positive/negative rates when the transmitter uses 6 and 24 subcarriers respectively. For SNR values greater than 15dB, the detection accuracy is above 95%.

Collision management and throughput: Latency constraints with the USRP platform disallow realtime evaluation. Therefore, we use a custom simulator to evaluate T2F against 802.11. We do not model the detailed properties of the wireless channel. We believe this is acceptable because our goal is to understand the collision probabilities in picking random subcarriers, and the improvements in channel utilization due to subcarrier based backoff. Thus, we operate at the granularity of time slots, and assume that packets fail only due to collisions. Figure 7(a) shows that T2F achieves low collision probability when using 52 subcarriers. It also shows that a single round of contention (in T2F-single-round) causes increasing collisions, justifying the necessity of a second round (802.11 is better than T2F-single-round because it increases its backoff exponentially to cope with congestion). Since *T2F-two-rounds* regulates the number of nodes promoted to the second round, the collision probability remains low and stable.

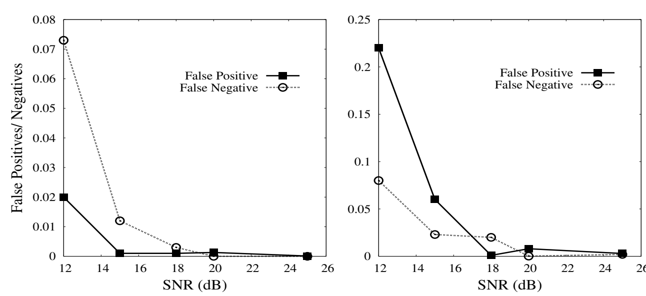

Figure 6: False positives/negatives for subcarrier detection, using (a) 6 subcarriers, 64pt IFFT (b) 24 subcarriers, 256pt IFFTs at the transmitter. For SNR values greater than 15dB, the detection accuracy is above 95%.

Figure 7(b) demonstrates the percentage throughput gain when the idle slots (due to 802.11 backoff) is better utilized with T2F. Evidently, the performance gap increases at higher rates because the idle slots occupy a relatively

[3] Recent work [6] has shown that this distance can be reduced significantly without increasing the strength of the self-signal – we have not adopted this optimization in this prototype.

larger portion of the channel time. Figure 7(b)) also shows a 10% throughput gain at 54Mbps from scheduling alone; we scheduled only 2 transmissions for this experiment.

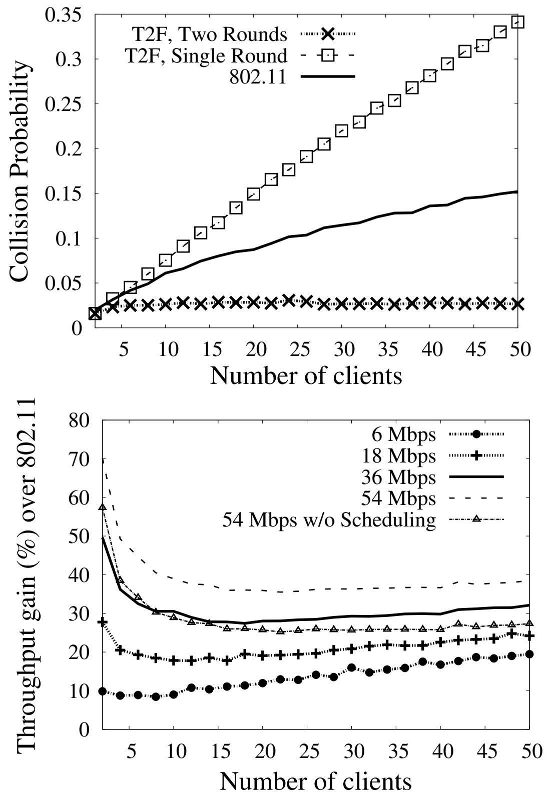

Figure 7: Performance of T2F vs 802.11: (a) collision probability. (b) relative throughput gain. Higher the rate better the gain. Scheduling (comparison shown only for the bitrate of 54 Mbps) alone gains around 10%.

5. RELATED AND ON-GOING WORK

Backoff-induced channel wastage in WiFi networks [1] has inspired lot of prior research. Proposals include adaptive backoff [2], implicit pipelining [8], intelligent queuing [1], etc. Very recently, [3] propose the possibility to signal on the frequency domain to enable fine grained frequency division multiplexing. While the ideas bear similarity to ours, they are not targeted towards contention resolution. [9] presents a scheme in which contention between active APs is resolved by other referee APs that provide feedback using OFDM subcarriers. Such an approach imposes tight synchronization requirements similar to [3]. On the other hand, T2F enables APs to perform local decision reducing backoff overhead and enabling TDMA-like schedule. To summarize, T2F breaks away from a long-standing method of contention resolution; to demonstrate success, it warrants a long term research agenda. This paper may be viewed as a first step toward this goal. Several questions remain open for ongoing work.

Robustness of subcarrier detection: The feasibility results in this paper are derived from laboratory experiments, without node/environment mobility. Subcarrier detection under harsh conditions needs to be tested extensively.

Coexistence with MIMO: T2F is complementary to MIMO or 802.11n systems, because an additional antenna in these systems may be utilized for contention resolution. In fact, the feasibility of higher data rates in these protocols emphasize the need to eliminate idle slots.

Interoperability with 802.11: Since T2F breaks away from convention, it is natural to ask whether nodes employing T2F can coexist with non-T2F nodes in an 802.11 network. Considering that T2F nodes still do carrier sense before transmitting data, we believe T2F can interoperate with 802.11 but possibly at the cost causing unfairness to non-T2F nodes. We are currently undertaking extensive analysis and evaluation of this interaction to study the feasibility and impact of incremental deployment of T2F.

6. CONCLUSION

Randomization is an effective method of contention resolution in systems with shared resources. Several protocols implement contention resolution by requiring nodes to wait for random durations. During this wait, the channel must remain idle, forcing undesirable under-utilization of channel. This paper proposes a nearly-instantaneous contention resolution method by observing the possibility to operate on the frequency domain (using OFDM subcarriers). A proof-of-concept on a small USRP testbed confirms feasibility and promising performance improvement. Developing a full-scale design is the natural next step.

7. REFERENCES

[1] A.P. Jardosh et al, “Understanding congestion in IEEE 802.11 b wireless networks,” in *ACM IMC*, 2005.

[2] Y. He, R. Yuan, J. Sun, and W. Gong;, “ Semi-Random Backoff: Towards resource reservation for channel access in wireless LANs,” *IEEE ICNP*, 2009.

[3] K. Tan et al, “Fine Grained Channel Access in Wireless LAN,” *ACM SIGCOMM*, 2010.

[4] V. Shrivastava et al, “Centaur: Realizing the full potential of centralized wlans through a hybrid data path,” in *MOBICOM*, 2009.

[5] A. Dutta et al, “SMACK: a SMart ACKnowledgment scheme for broadcast messages in wireless networks,” *ACM SIGCOMM*, 2009.

[6] J. Choi et al, “Achieving single channel, full duplex wireless communication,” *MOBICOM*, 2010.

[7] M. Vutukuru et al, “Cross-Layer Wireless Bit Rate Adaptation,” in *SIGCOMM*, 2009.

[8] X. Yang and N.H. Vaidya, “A wireless MAC protocol using implicit pipelining,” *IEEE Transactions on Mobile Computing*, 2006.

[9] B. Roman et al, “Multi-carrier burst contention (MCBC): Scalable medium access control for wireless networks,” in *IEEE WCNC*, 2008.