Listen Before You Talk
But On The Frequency Domain

Souvik Sen   Romit Roy Choudhury   Srihari Nelakuditi
Duke University   Duke University   University of South Carolina

Abstract

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 device that reaches zero first wins the contention and initiates transmission. This core idea – called backoff – is known to be inherently wasteful because the channel must remain idle while all contending nodes are simultaneously counting down. This wastage has almost been accepted as a price for decentralization.

In this paper, we ask whether the entire backoff operation can be migrated to the frequency domain, thereby eliminating a long-standing source of channel inefficiency. Our core idea draws on OFDM subcarriers in modern WiFi radios, treating each subcarrier as an integer number. By transmitting a signal on a random subcarrier frequency, and using a second antenna to listen for all active frequencies, each node may be able to detect its *rank* among all contenders. Since signaling on subcarriers is almost instantaneous, the wastage from backoff can become negligible. We design such an unconventional backoff scheme called *Back2F*, implement it on a software-radio testbed, and demonstrate its feasibility with real-world experiments. A natural next step would be to revisit today's protocols, and ask what other operations may be similarly migrated to the frequency domain.

I. Problem & Motivation

Accessing a shared resource in a decentralized manner has been studied for several decades. In the networking context, ALOHA systems showed how random access strategies can be effective when multiple computers communicate over a shared channel. The basic intuition suggests that each device pick a random number from a certain range and begin counting down from it – the device that reaches zero first wins this contention and begins transmission. The losers wait for this transmission to finish, and repeat the random count-down procedure, thereafter. Since every node counts down at the same pace, 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. This core idea – called backoff – has been optimized over time, and underpins several modern day technologies, including Ethernet [1], Inmarsat satellite systems [2], WiFi, etc. While backoff is indeed simple and effective, it is known to be wasteful because the channel must remain idle while all nodes perform their count-down in time. Moreover, under high network density, multiple nodes are likely to chose the same random backoff, resulting in frequent packet collisions. Packet collisions trigger an exponential increase in the backoff value, resulting in greater channel wastage.

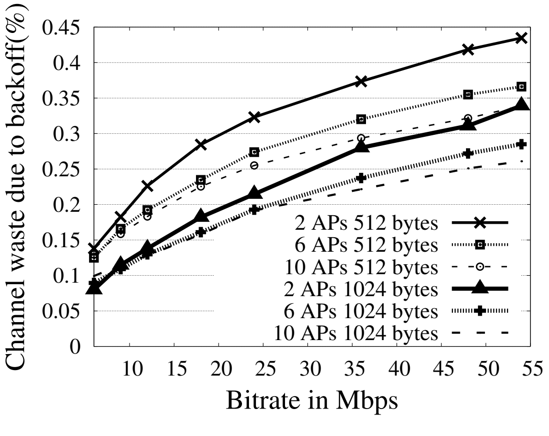

Fig. 1. Overhead of 802.11 backoff. Larger fraction of channel wasted with smaller packets at higher bitrates.

We make three observations. (1) Fundamentally, backing off is not a time domain operation. Its conventional implementation has been in the time domain, forcing the channel to be idle before each packet transmission. (2) The duration of each backoff timeslot is fixed, so the channel wastage is relatively higher with smaller packets and higher bitrates. 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. Also, the slot duration has fundamental limits, lower bounded by the propagation delay, clear channel assessment (CCA) time, and circuit delays [5]. Hence, a slot cannot become much shorter than the current $9\mu s$. (3) Finally, although channel utilization may improve with few nodes (backing off in parallel), just a few more nodes can cause collisions (802.11 experiences 18% collisions with 3 nodes choosing their backoff values from the range $[0, CW]$, where $CW = 16$ is called the contention window). A collision forces nodes to exponentially increase their CW, pushing the system back to under-utilization. Fig. 1 shows the channel under-utilization due to 802.11's backoff, under varying bitrates and network densities. Authors in [3], [5] corroborate these findings with extensive analysis and measurements, emphasizing the need to improve wireless contention resolution. This paper attempts to address this problem by migrating the backoff operation to the frequency domain.

II. RELATED WORK

The notion of backoff dates back to 1973, when pure/slotted ALOHA systems [6] were originally proposed (see [7] for a history on spectrum sharing). The core ideas from ALOHANet have found wide applicability in Ethernet, the Inmarsat satellite network, and most recently, in WiFi [8], [9]. With WiFi's popularity, binary exponential backoff became a heavily researched topic. One thread of proposals have optimized the manner in which backoff adapts to collisions and network conditions [8], [10], [11]. While these proposals are appealing for their simplicity, practical measurements [3], [12] and analytical studies [13] show that the inherent inefficiencies remain, and become pronounced in unfavorable conditions.

In another research thread, researchers attempted to adapt the backoff scheme based on estimations of network traffic/contention [14]. Similar proposals reduces contention overhead using TDMA-style scheme (TCF [15]), centralized control ([16]). Unfortunately, such schemes are prone to error, due to unpredictable variation in traffic and interference patterns [17]. Also, they do not scale to chaotic networks. FICA [5] showed the possibility of signaling on the frequency domain to facilitate fine grained FDMA. FICA requires involved RTS/CTS exchanges and a common "referee" node to perform the arbitration, similar to ideas in [18]. Also, the approach in [18] relies on tight time synchronization, that may experience practical challenges in a real system.

III. A UNIQUE APPROACH: BACKOFF IN FREQUENCY DOMAIN

We observe that backoff is fundamentally a method to rank nodes in a decentralized manner, and there is no reason that it has to be executed in the time domain. If multiple nodes can mutually share their randomly chosen values, then each node can learn about the global ranking – the winner will be obvious to all. Our contribution lies in recognizing that modern OFDM systems offer an opportunity to realize "distributed ranking" in the frequency domain. In OFDM, a packet is transmitted by striping bits over multiple narrow band frequency subcarriers. The 802.11a/g 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. In our protocol, each node announces its random backoff by signaling on the corresponding subcarrier, while a second antenna on the same node listens to all the active subcarriers. By observing the relative position of its own subcarrier against all other active subcarriers, a node can determine if its is the winner. With this background we present `Back2F` in 3 sub-parts:

A. Backoff within a Single Collision Domain

Figure 2 shows an example where contenders AP1 and AP2 observe the same channel, and hence, when the channel becomes idle, both APs recognize it as a trigger to begin transmission. They choose random numbers 11 and 29, respectively. However, instead of counting down these numbers in time, they transmit a short signal on their corresponding subcarriers.

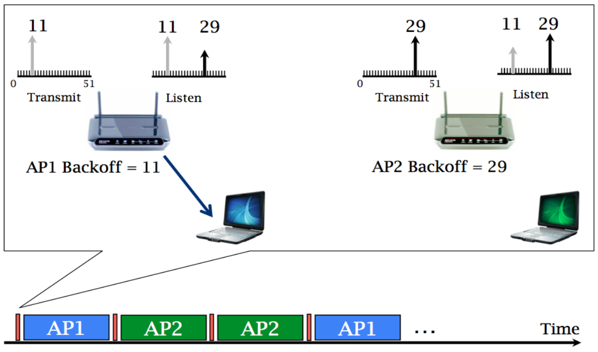

Fig. 2. A close up view of the first backoff. AP1 picks/activates subcarrier 11 and AP2 chooses 29. They learn of other backoff values through subcarriers. AP1 with smaller backoff transmits whereas AP2 defers.

While the signal is being transmitted on the transmit antenna, a 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 then extracts all the active subcarriers, thereby learning the backoff values of the other contenders. With knowledge of everyone's backoffs, each AP can instantaneously determine whether it has won the contention.

AP1 with smaller backoff of 11 proceeds to transmit, whereas AP2 with larger backoff defers. Of course, data transmission is performed using all the subcarriers, identical to regular 802.11. The deferred node deducts the smallest known backoff and contends again after the channel has become idle, i.e., after AP1 finishes, AP2 contends with a backoff value of $29-11=18$. Observe that the net effect is exactly like time-domain backoff in 802.11. All the contending nodes count down simultaneously till the smallest of them reaches zero; the node whose backoff reaches zero proceeds to transmit, while others contend later with their reduced backoff values. However, unlike 802.11, the time to pick a winner in `Back2F` is much shorter, in the order of few OFDM symbols.

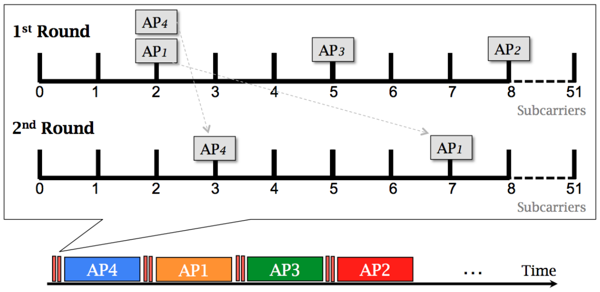

Fig. 3. Illustration of Back2F with two contention rounds. AP1 and AP2 choose the same smallest backoff, and enter the second round of contention. AP4 wins the second round and accesses the channel.

What if two contending nodes choose the same backoff value – how does `Back2F` cope with collisions? Collisions are certainly possible when two nodes pick the same random subcarrier. `Back2F` copes with collisions by introducing a second round of subcarrier based contention. A node that is a winner in the first round, retransmits on another randomly chosen subcarrier immediately after. Figure 3 illustrates the process. AP1 and AP4 choose the same backoff value that happens to be smallest among all the other backoffs. Then, both APs advance into a second round of contention, and this time AP4 picks 3 while AP1 picks 7. AP4 being the winner of the second round proceeds to transmit while AP1 waits to participate in the next backoff.

B. Backoff over Multiple Collision Domains

For the ease of introduction, the above description of `Back2F` makes a simplifying assumption that all the contending nodes can hear each other. Thus, they become aware of all backoff values, resulting in a consistent view of the global ranking among nodes. In practice, however, a wireless network will obviously span over multiple collision domains as in Fig. 4. In this example, AP1 and AP2 belong to one collision domain, while AP2, AP3 and AP4 belong to a different one (i.e., AP1 does not carrier sense AP3 or AP4, and the vice versa). When a node such as AP2 belongs to two collision domains, it may be a winner in one but not in the other. `Back2F` copes with these cases as described below.

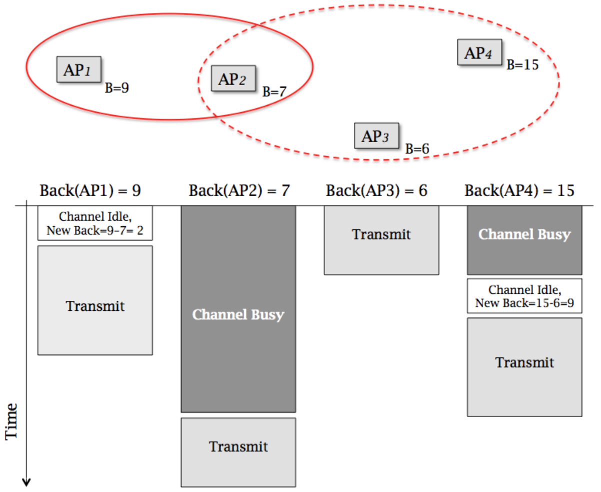

Fig. 4. `Back2F` with multiple collision domains: Due to differing views, AP1 thinks AP2 won the contention whereas AP2 lost to AP3. However, when the channel is idle for DIFS, AP1 performs backoff and transmits.

Suppose the backoff values of AP1, AP2, AP3, and AP4 are 9, 7, 6, and 15, respectively. Then, according to AP1, node AP2 is the winner, whereas in AP2's view, the winner is AP3 The consequence is that only AP3 proceeds to transmit, AP2 defers to AP3, and AP1 defers to AP2. This is unnecessary because AP1 could very well transmit in parallel to AP3. `Back2F` addresses this form of *head-of-line blocking* to uphold spatial reuse in the network. When AP1 observes that the channel is idle for DIFS duration, it infers that the winner is blocked by some other transmission. Hence, AP1 initiates a backoff with its revised value of $9-7=2$. Assuming AP1 is the only contender, it wins the channel and begins transmitting. Now, even though AP3 completes transmission, AP2 still does not transmit because it carrier senses AP1. AP4 now observes an idle channel, readjusts its backoff to $15-6=9$, and advances into communication. Once AP1 and AP4 are done, AP2 transmits its packet. Again, observe that the overall order of transmissions mimics 802.11, only the backoff procedures are quicker.

C. Optimization: Batched Transmissions

Unlike 802.11, in `Back2F` each node knows the exact backoff values chosen by other nodes This enables each node to learn its rank *a priori* in the sequence of pending transmissions. `Back2F` exploits this knowledge to batch transmissions (i.e., a train of back-to-back packets between successive backoffs).

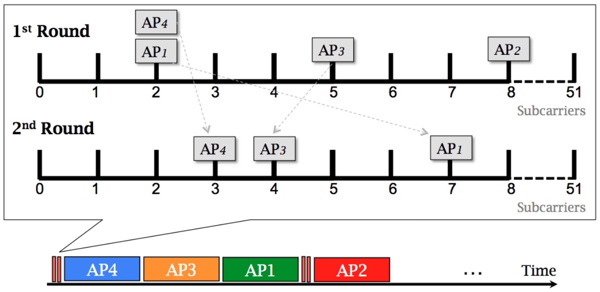

Fig. 5. Backoff in the frequency domain followed by scheduled transmissions. All APs contend in the first round, but only AP1 and AP2 enter the second round. Based on backoff values in the second round, the schedule is AP1 followed by AP2. Only after the scheduled transmissions complete, AP3 and AP4 contend again.

As a first step, `Back2F` promotes the *top-K* ranked APs to the second round, as opposed to the winner(s) alone. Higher values of K will result in longer batch size (better throughput), but at the risk of collisions. We set $K = 3$ in our implementation, empirically observed to provide the best tradeoff. By observing the active subcarriers in the first round, each node can independently decide whether to participate in the second round. As an example, suppose 4 APs in Figure 5 contend in the first round, and the protocol intends to promote the *top-2* ranked APs to the second round. AP1 and AP4, both ranked 1, naturally advance; AP3 also advances because it is ranked 2. These three nodes are then ordered as AP4, AP3, AP1 in their second round of contention, and they transmit back-to-back in that sequence, *without a per-packet backoff.*

IV. IMPLEMENTATION AND EVALUATION

This section is organized to answer three main questions: (1) the reliability of subcarrier detection in an actual prototype network, (2) `Back2F`'s collision probability with increasing density, and (2) `Back2F`'s performance gain in realistic network topologies. We begin the discussion with a description of our prototype testbed.

USRP/GNURadio Prototype: We prototype `Back2F` on a small testbed of 10 USRPs. Both the transmitter and receiver uses a 8MHz band. The transmitter is equipped to transmit on any of the 52 subcarriers, that are converted into a time domain signal using 64pt IFFT. The listening antenna executes 256 point FFT to offset the effect of high self signal from the transmitter. The listener antenna detects subcarriers using a joint thresholding and peak-detection scheme. Whenever a peak is above a threshold, `Back2F` declares it as an active subcarrier. Since backoff is always preceded by a DIFS interval in which the channel is idle, this threshold is adaptively chosen by sampling the noise and interference floor. This helps in keeping the false positives/negatives low.

Subcarrier Detection: The feasibility of detecting a subcarrier, in presence of a strong self-signal, is the problem of interest. Figure 6(a) shows the feasibility of discerning multiple adjacent subcarriers using 256pt. FFT. The subcarrier with the highest spike is the self-subcarrier, while contending nodes were made to transmit on adjacent and nearby subcarriers. A higher point FFT reduces this leakage; even in presence of a high self-subcarrier, transmissions on adjacent and nearby subcarriers are reliably discerned. Figure 6(b) shows the detection accuracy $(1-$ FalseNegative$)$ as a function of subcarrier distance from the self-subcarrier. As anticipated, the influence of the self-signal reduces with increasing distance. Using a 256 pt FFT, subcarriers above 14dB can be detected reliably. False positives in our experiments were rare at less than 2%. Carrier sense threshold in 802.11g/n permits transmission when the signal in the channel is 13dB or below in comparison to the noise floor [19], [20]; thus `Back2F` will almost be able to detect all links that are within the collision domain.

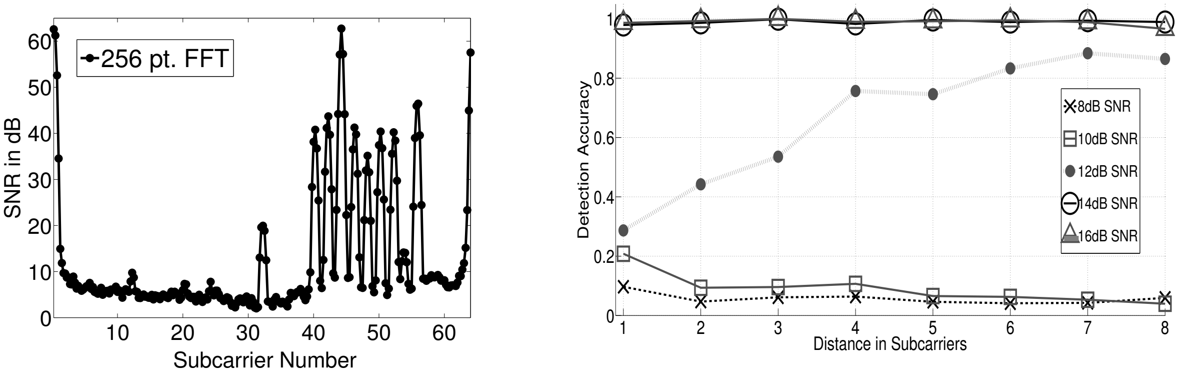

Fig. 6. (a)Active subcarrier detection when 10 nodes transmitting with 64pt IFFT using 8MHz bandwidth. Receiver using 256 pt. FFT. (b) Detection accuracy of subcarriers at varying distance from the self subcarrier with 256pt FFT at receiver: Adjacent subcarriers with SNR 14dB or greater can be detected with 97% accuracy.

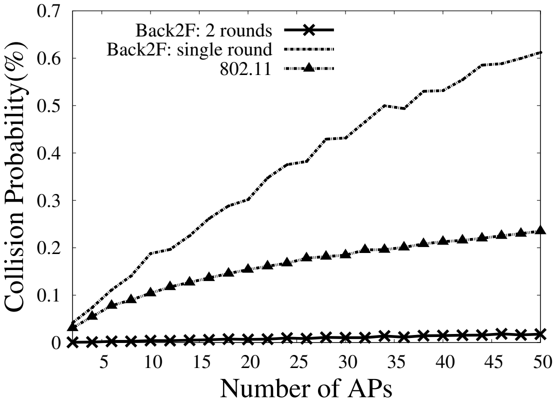

Fig. 7. Collision probability of `Back2F` with two rounds remains below 2% in high density networks with 52 subcarriers. 802.11 experiences more collisions.

Collision probability: Figure 7 shows the collision probability with increasing number of contenders. Observe that when using 52 subcarriers, the collision probability increases quickly with a single round of contention, and in fact, is worse than 802.11. However, having a second round in `Back2F` drastically reduces the collision probability. Evidently, even when the number of contenders is more than 50, the collision probability still remains less than 2%. 802.11, on the other hand, collides far more frequently, and adapts through exponentially increasing its contention window. The adaptation is an heuristic, and naturally incurs a performance penalty.

Throughput Evaluation: The above USRP/GnuRadio based prototype is suitable for demonstrating the feasibility of active subcarrier detection, but not the resulting gain from `Back2F`. Latency constraints with the USRP platform disallow realtime evaluation. Therefore, we resort to trace based evaluation to assess the efficacy of Back2F. We collected RSSI, channel response, bitrate, interference map related traces to mimic realistic traffic and topologies. Figure 8(a) reports the throughput gain with `Back2F` for Skype, Web browsing, and HD streaming traffic. Evidently, the benefits of `Back2F` are available across all these classes of traffic. Unsurprisingly, gains are better with Skype traffic due to smaller packet sizes. This is because backoff overheads are fixed, making it proportionally larger for to short packet transmissions. Figure 8(b) shows `Back2F`'s throughput gain over 802.11, for varying bitrates and node density. It also presents the performance of `Back2F` without batch at 54Mbps. Across all settings, `Back2F` provides gains are in the range of 10% to 55%, suggesting the possibility to scale to large networks.

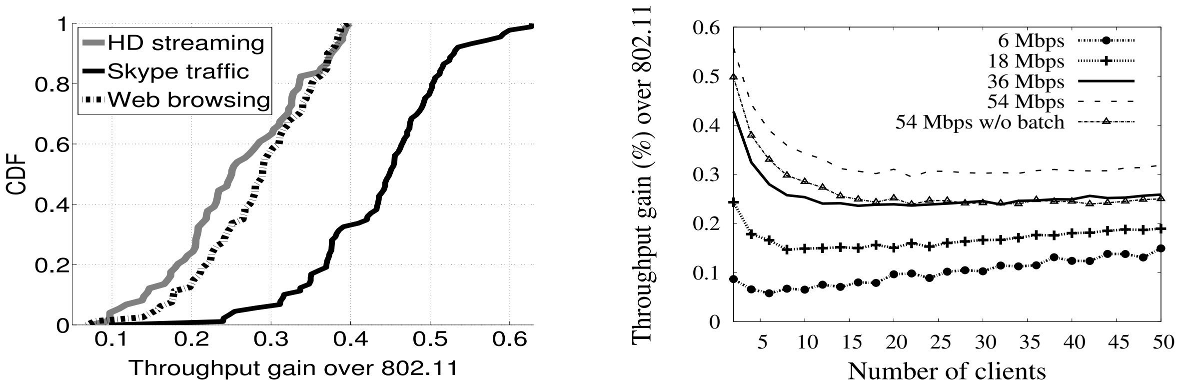

Fig. 8. (a) Different types of traffic: Gain with `Back2F` is more pronounced for Skype traffic with smaller packets. (b) Performance of `Back2F` in single collision domain: Higher the rate better the gain. Batching (comparison shown only for 54 Mbps) offers around 6% gain.

V. CONCLUSION AND CONTRIBUTION

Randomization is an effective method of contention resolution in systems with shared resources. Several media access 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 leveraging the possibility to operate on the frequency domain (using OFDM subcarriers). Our contributions in this paper may be summarized as:

- *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 arbitration 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 larger system.

More importantly, this paper stimulates the community to approach system design from a new perspective. It shows the benefit of migrating time domain operations to frequency domain and provokes one to ask a more important question: "what other time domain operations can potentially be migrated to frequency".

REFERENCES

[1] R.M. Metcalfe and D.R. Boggs, "Ethernet: Distributed packet switching for local computer networks," *Communications of the ACM*, 1976.

[2] Inmarsat satellite network, "http:/http://www.inmarsat.com/," .

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

[4] Y. He and et al, " Semi-Random Backoff: Towards resource reservation for channel access in wireless LANs," *IEEE ICNP*, 2009.

[5] K. Tan et al, "Fine Grained Channel Access in Wireless LAN," *ACM SIGCOMM*, 2010.

[6] N. Abramson, "The Aloha System: another alternative for computer communications," in *Proceedings of the Fall joint computer conference*. ACM, 1970.

[7] C. Jackson, "Dynamic sharing of radio spectrum: A brief history," in *IEEE DySPAN*, 2005.

[8] V. Bharghavan, A. J. Demers, S. Shenker, , and L. Zhang, "Macaw: A media access protocol for wireless lan's," in *SIGCOMM*, 1994.

[9] IEEE Computer Society, "802.11: Wireless lan medium access control (mac) and physical layer (phy) specifications," 1997.

[10] T. Nandagopal, T.E. Kim, X. Gao, and V. Bharghavan, "Achieving MAC layer fairness in wireless packet networks," in *ACM MOBICOM*, 2000, pp. 87–98.

[11] Q. Pang and et. al., "Performance evaluation of an adaptive backoff scheme for WLAN," *Wireless Communications and Mobile Computing*, vol. 4, no. 8, pp. 867–879, 2004.

[12] A.P. Jardosh et al, "IQU: practical queue-based user association management for WLANs," in *MOBICOM*, 2006.

[13] G. Bianchi, "Performance analysis of the IEEE 802.11 distributed coordination function," *IEEE Journal on selected areas in communications*, vol. 18, no. 3, pp. 535–547, 2000.

[14] Z.J. Haas and J. Deng, "On optimizing the backoff interval for random access schemes," *Communications, IEEE Transactions on*, vol. 51, no. 12, pp. 2081–2090, 2004.

[15] C. Lim and C.H. Choi, "TDM-based coordination function (TCF) in WLAN for high throughput," in *GLOBECOM*, 2005.

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

[17] G. Bianchi and I. Tinnirello, "Kalman filter estimation of the number of competing terminals in an IEEE 802.11 network," in *INFOCOM*, 2003.

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

[19] M.Z. Brodsky and R.T. Morris, "In defense of wireless carrier sense," in *ACM SIGCOMM*, 2009.

[20] "Ieee 802.11-2007 standard," *standards.ieee.org/getieee802/download/802.11-2007.pdf*.