

requirement is specific to USRPs since the current 802.11a/g OFDM designs on USRPs need higher FFT sizes due to imprecision [26, 31]. We believe commercial hardware offers greater precision, obviating the need for higher FFT sizes at the listener.

The listening antenna detects subcarriers using a joint thresholding and peak-detection scheme. This is necessary because with practical hardware (especially USRPs), the subcarriers emerge as peaks rather than ideal impulses (Figure 8). 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 over this interval. 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. In the test, transmitters were randomly placed and made to transmit signals on subcarriers at varying spectral separation from the self-subcarrier. Figure 11 shows the detection accuracy ($1 - \text{FalseNegative}$) as a function of subcarrier distance from the self-subcarrier. As anticipated, the influence of the self-signal reduces with increasing distance. Also, with increasing FFT size at the listening antenna, even the nearby subcarriers can be detected more accurately. Using a 256 pt FFT, subcarriers above 14dB can be detected reliably. OFDM based 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 [1, 8]; thus Back2F will almost be able to detect all links that are within the collision domain. Occasional false negatives may still happen, however, as we show later, Back2F is reasonably robust to such occurrences.

Impact of noise and interference: Back2F is most vulnerable to misdetection on the subcarrier adjacent to the self-subcarrier. Hence, we focus on the detection of this subcarrier to investigate the worst case performance of Back2F. Figure 12(a) shows the false negatives in detecting the adjacent subcarrier with 256pt FFT, for varying SNR of the signal on that subcarrier. It also shows the false positives, i.e., incorrectly detecting an inactive subcarrier. Clearly, false positives are rare, less than 2%. As the SNR of the signal on the adjacent subcarrier increases, the possibility of false negatives decreases, and at 14dB or more, it can be detected reliably.

Figure 12(b) also shows that background interference does not affect our results. Back2F does not falsely detect background interference as an active subcarrier because it adjusts its threshold (as explained above) depending on energy per-subcarrier estimated during the idle DIFS. We believe that at higher bandwidths (20MHz in 802.11), Back2F’s detection accuracy may improve as a consequence of greater spectral separation between subcarriers. Due to limitations in the processing capability of USRPs, the experiments in this paper are performed at 8MHz bandwidth, but with all 52 subcarriers as in 802.11. Moreover, results presented later show that Back2F can tolerate some incidence of false negatives.

4.2 Trace based Performance Evaluation

The above USRP/GNUradio based prototype is suitable for demonstrating the feasibility of active subcarrier detection,

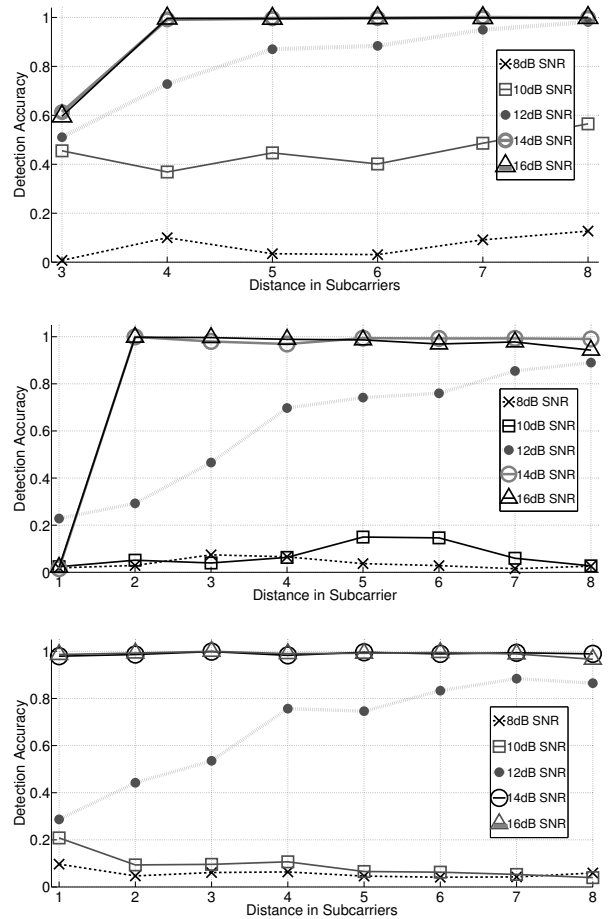


Figure 11: Detection accuracy of subcarriers at varying distance from the self subcarrier: Transmitter uses 64pt IFFT while listening antenna employs (a) 64, (b) 128, or (c) 256 pt FFT. With 256pt FFT, adjacent subcarriers with SNR 14dB or greater can be detected with 97% accuracy. Similar results for 128pt and 64pt FFT but at a higher separation. With 128pt and 64pt FFT, we can use every 2nd and 4th subcarrier respectively for Back2F.

but not the resulting gain from Back2F. Latency constraints with the USRP platform disallow realtime implementation of all the Back2F protocol operations. Therefore, we resort to trace based evaluation to assess performance in realistic scenarios.

Evaluation setting: To conduct high fidelity emulation of real world setting, we collected traces of channel characteristics and network traffic as follows. We placed APs in 20 locations and clients in 45 locations throughout our engineering building. To gather information per subcarrier, we use Intel 5300 chipset based wireless cards [2]. For each AP to client link, we recorded the RSSI, channel impulse response, transmission bitrate, and collision probability with respect to the strongest interfering AP. The transmission bitrate is experimentally selected as the highest bitrate that can support a delivery ratio of 90% or more. To estimate collision probability, we turn off carrier sensing at the APs and activate downlink transmissions in pairs. The collision probability is experimentally calculated with 10 runs of 500 packets of size 1500 bytes. We also gath-




Figure 12: False positives/negatives in subcarrier detection with (a) 256pt FFT at receiver; (b) 256pt FFT in presence of 10dB interference. Low false positives in general and false negatives are low too at 14dB or more. Also, subcarrier detection is tolerant to interference.

ered RSSI and channel impulse response between every pair of APs. The Intel 5300 cards combine impulse responses to report the channel matrix for only 30 subcarriers. We estimate the channel matrix for all the 52 subcarriers via interpolation. As a representative of traffic mix in the real world, we collected traces from Skype (real time traffic), web browsing and HD streaming sessions. The average packet size in these cases were 511, 1063, and 1424 bytes respectively.

Based on the collected traces, we emulated topologies of various sizes. To model a topology with k transmitters, we uniformly choose k APs to cover the building. The remaining $65-k$ nodes are treated as clients and each of them is associated with the nearest AP (with the strongest RSSI) creating a wireless LAN like setting. We pick 100 instances of each topology size ranging from 6 APs to 18 APs. We use the real world traffic traces collected above to emulate download traffic from APs to clients. When an AP has a packet to transmit, we use the traces to determine which other APs can carrier sense this transmission, the collision probability with another hidden AP, its ability to detect active subcarriers, etc. These attempts are targeted to mimic real-world scenarios.

The relative performance of Back2F over 802.11 depends on the transmission bitrates — higher the bitrate, better the relative gain with Back2F. Hence, it is important to report the nature of links in the emulated topologies in our evaluation. Figure 13 shows the CDF of bitrates of links in each topological setting with varying number of APs. It indicates that our topologies included links with several different bitrates, not just the highest bitrate to favor Back2F.

Throughput gain: We compare Back2F’s overall throughput against 802.11, under varying topologies and Internet traffic patterns. Figure 14(a) and (b) present the relative throughput gain due to Back2F without and with batched transmissions, respectively. The number of APs in the topologies vary from 6 to 18, while clients range from 59 to 47. Back2F consistently outperforms 802.11 across all scenarios. Batched




Figure 13: Bitrates of links in the emulated topologies.

transmissions further reduce backoff overhead and improve throughput by around 5%. The gains with batching is small because basic Back2F has already reduced the backoff duration substantially.




Figure 14: Throughput gain with Back2F over 802.11: (a) without and (b) with batching of transmissions by Back2F. Throughput gain with Back2F ranges from 5% to 38%. Batching contributes 5% of the gain.

Traffic type: Figure 15 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 short packets.

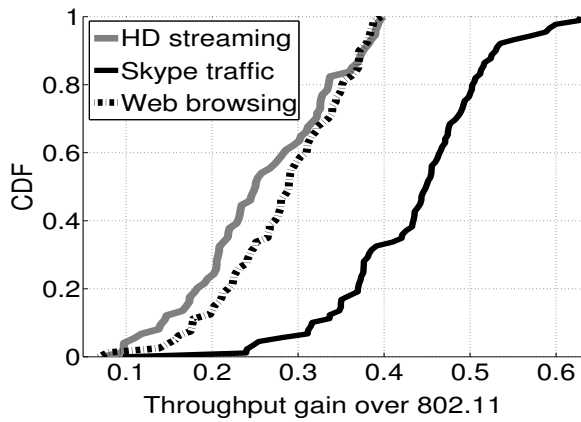


Figure 15: Different types of traffic: Gain with Back2F is more pronounced for Skype traffic with smaller packets.

Fairness: Since Back2F emulates the countdown of 802.11, it is expected to be similar in fairness to 802.11. To verify that, we compute Jain’s fairness index on throughput obtained by each AP. Figure 16 shows the fairness index with Back2F and 802.11 for different topology settings. It gives the mean index and confidence interval over 100 instances for each topology size. Clearly, Back2F offers throughput gains, while sustaining fairness comparable to 802.11.

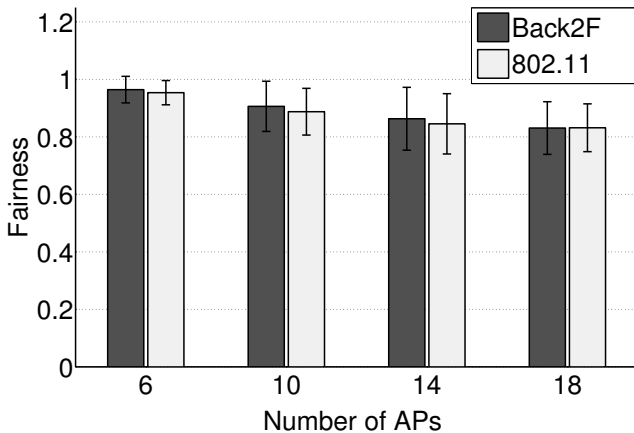


Figure 16: Fairness: Back2F emulates countdown of slots in 802.11 and provides similar fairness.

Impact of channel fading: A Back2F node may fail to detect an active subcarrier (false-negative). We study the impact of such misdetection on Back2F. To cope with misdetection, as explained in Section 3.5, in the second round, a node that picks random number i , transmits on subcarriers i and $((26 + i) \bmod 52)$. Figures 17 show the degradation from ideal throughput due to false-negatives (ideal throughput obtained with no false positives/negatives). Even with 20% of false-negatives, the resulting throughput degrades by only 5%. In essence, Back2F is a viable scheme that is not overly sensitive to subcarrier misdetection.

Dense networks: The above evaluation investigates the performance of Back2F with up to 18 APs placed in our engineer-

ing building. To investigate its scalability to denser networks, we simulated HD traffic in a single collision domain under varying densities and different bitrates. Figure 18 shows Back2F’s throughput gain over 802.11, when up to 3 transmissions were allowed per batch. It also presents the performance of Back2F without batch at 54Mbps. Across all settings, Back2F provides gains are in the range of 15% to 30%, suggesting the possibility to scale to large networks.

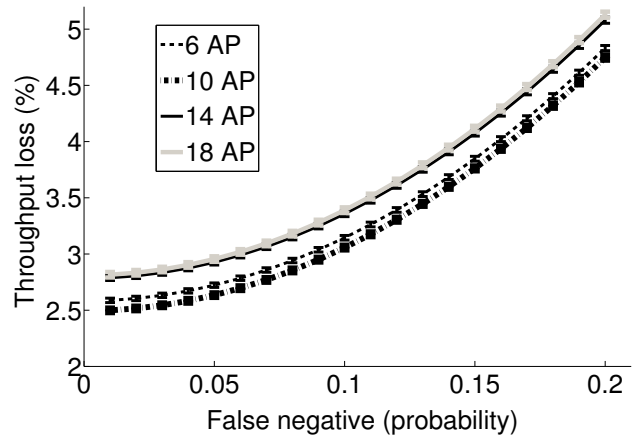


Figure 17: Impact of false negatives on Back2F throughput is shown in the form of deviation from ideal throughput. Due to redundant activation of subcarriers in the second round, the effect of false negatives is minimal.




Figure 18: Performance of Back2F in single collision domain: Higher the rate better the gain. Batching (comparison shown only for 54 Mbps) offers around 6% gain.

5. LIMITATIONS AND ON-GOING WORK

Back2F breaks away from a long-standing method of contention resolution; to demonstrate success, it warrants continued research engagement. This paper may be viewed as a first step toward this goal. Several extensions and enhancements remain open for future work.

Robustness of subcarrier detection: The feasibility results in this paper are derived from lab experiments, without node/

environment mobility. Compared to time-domain backoff, Back2F may be more sensitive to channel fluctuations. As discussed earlier, subcarrier detection can be made more robust by stripping signals over multiple subcarriers, in order to convey a backoff value. We need to investigate such techniques further and carefully evaluate subcarrier detection under harsh conditions.

Collisions due to hidden terminals: Contention resolution schemes are not designed to cope with hidden terminal problems. However, when using 802.11, the exponential increase in backoff may eventually separate the hidden terminals in time, permitting a successful transmission. Of course, once a success occurs, 802.11 resets its contention window, bringing back the hidden terminal problem. With Back2F, the hidden terminals would continue to collide if they continue to transmit, and in that sense, 802.11 might be slightly better. Nevertheless, we observe that Back2F collisions are solely due to hidden terminals; collisions caused by identical backoff values are far less likely in Back2F. Thus, Back2F can confidently diagnose the cause of collisions, and perhaps turn on RTS/CTS in the face of collisions. 802.11, on the other hand, would still need to discriminate between the cases of identical-backoff and hidden terminals.

Need for an additional antenna: Back2F has to transmit and listen simultaneously only during backoff. The listening antenna can very well act as an additional antenna during normal transmission/reception, such as in a MIMO system. In other words, Back2F is complementary to MIMO. In fact, the feasibility of higher data rates with MIMO emphasizes the need to eliminate idle slots, and thereby adopt Back2F-like schemes. Even without MIMO, given that there are other uses of an additional antenna [25], its inclusion in WiFi devices may very well be worthwhile.

Gain over packet aggregation: 802.11n uses packet aggregation to reduce the contention overhead. The natural question then is whether Back2F is still beneficial. Depending on the type of traffic (e.g., VoIP), aggregation may not be possible nor suitable [27, 29, 30]. Even with packet aggregation, Back2F provides gains at high bitrates. Besides, Back2F addresses a fundamental problem of resolving contention keeping the channel utilization high, regardless of the traffic pattern.

Interoperability with 802.11: We believe Back2F can interoperate with (legacy) 802.11 nodes but may cause unfairness to them. A potential approach to alleviate unfairness is to have Back2F wait for longer than DIFS before participating in a backoff. This gives legacy nodes opportunity to countdown and eventually transmit. We need to understand this interaction, and study the feasibility of incremental deployment of Back2F.

Analysis and Correctness: Back2F emulates the countdown of 802.11 and therefore we believe it is similar to 802.11 in correctness and fairness. We have simulated Back2F for more than 48 hours on various network topologies – we have not encountered deadlocks, starvations, or other correctness problems. However, we have not formally analyzed Back2F's correctness properties. We leave an analytical treatment of Back2F to future work.

6. RELATED WORK

The notion of backoff dates back to 1973, when pure/slotted ALOHA systems [4] were introduced (see [14] 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 [3, 5]. With WiFi's popularity, exponential backoff became a heavily researched topic. Discussing this entire literature is difficult – we only discuss representative ideas, and discriminate them from Back2F.

Regulating Increase/Decrease: One thread of proposals have optimized the manner in which backoff adapts to collisions and network conditions. MACAW [5] proposes doubling of the backoff upon packet loss, but decrease of 1 upon success. PFCR proposed similar policies, but from the fairness perspective [19]. While these and other schemes [20] were appealing for their simplicity, practical measurements [15, 16] and analytical studies [6] show that the inherent inefficiencies remain, and become pronounced in unfavorable conditions.

Contention Estimation: In another research thread, researchers attempted to adapt the backoff scheme based on estimations of network traffic/contention [12]. Unfortunately, such estimations are not always reliable due to unpredictable variations in traffic patterns [7].

Scheduling (Centralized and Distributed): TCF [17] eliminates contention overhead by allocating the channel dynamically using a TDMA-style scheme. Noting the difficulties with synchronization and prediction in TDMA, ZMAC [22] proposed a hybrid MAC allowing CSMA for low contention environments and TDMA for high contention regimes. While creative, performance degradation in low contention regimes, as well as heavy coordination overhead, makes ZMAC impractical for dense networks. Several centralized solutions leverage a central controller to schedule transmissions in single-administrator environments, like offices, airports, etc. [18, 28]. Unfortunately, they do not scale to chaotic networks, such as in residential WLANs or for MiFi networks among personal devices.

PHY based Techniques: Recently, PHY capabilities are being leveraged to redesign higher layer protocols [10, 11, 21, 25, 29]. FICA [29] showed the possibility of signaling on the frequency domain to facilitate fine grained FDMA. While some ideas bear resemblance to Back2F, FICA requires involved RTS/CTS exchanges and a common “referee” node to perform the arbitration, similar to ideas in [23]. Also, the approach in [23] relies on tight time synchronization, that may experience practical challenges in a real system. The protocol in [24] performs frequency domain backoff, however, it does not address multiple collision domains, fairness, starvation, fading, etc. Evaluation is also limited. Back2F demonstrates the feasibility of a distributed backoff mechanism, free of RTS/CTS and referee nodes. The use of an additional antenna, a second round of contention, resilience to fading, packet batching, and a functional prototype, discriminate Back2F from existing work. We believe that Back2F is a departure from established ideas in time-domain contention resolution, begging the question: *what else can be migrated to the frequency domain?*

7. 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 exploiting the opportunity 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.

8. ACKNOWLEDGEMENT

We sincerely thank the anonymous reviewers for their valuable feedback on this paper. We are also grateful to NSF for partially funding this research through the following grants – CNS-0448272, CNS-0917020, CNS-0916995, and CNS-0747206.

9. REFERENCES

- [1] IEEE 802.11-2007 standard. *standards.ieee.org/getieee802/download/802.11-2007.pdf*.
- [2] Intel802.11n channel measurement tool. *ils.intel-research.net/80211n-channel-measurement-tool*.
- [3] 802.11: Wireless LAN medium access control (MAC) and physical layer (PHY) specifications, 1997.
- [4] ABRAMSON, N. The Aloha System: another alternative for computer communications. In *Proceedings of the Fall joint computer conference (1970)*, ACM.
- [5] BHARGHAVAN, V., DEMERS, A. J., SHENKER, S., , AND ZHANG, L. MACAW: A media access protocol for wireless lan's. In *SIGCOMM (1994)*.
- [6] BIANCHI, G. Performance analysis of the IEEE 802.11 distributed coordination function. *IEEE Journal on selected areas in communications* 18, 3 (2000), 535–547.
- [7] BIANCHI, G., AND TINNIRELLO, I. Kalman filter estimation of the number of competing terminals in an IEEE 802.11 network. In *INFOCOM (2003)*.
- [8] BRODSKY, M., AND MORRIS, R. In defense of wireless carrier sense. In *ACM SIGCOMM (2009)*.
- [9] CHOI, J., JAIN, M., SRINIVASAN, K., LEVIS, P., AND KATTI, S. Achieving single channel, full duplex wireless communication. *ACM MOBICOM (2010)*.
- [10] DUTTA, A., SAHA, D., GRUNWALD, D., AND SICKER, D. SMACK: a SMart ACKnowledgment scheme for broadcast messages in wireless networks. *ACM SIGCOMM (2009)*.
- [11] GOLLAKOTA, S., AND KATABI, D. Zigzag decoding: Combating hidden terminals in wireless networks. *ACM SIGCOMM (2008)*.
- [12] HAAS, Z., AND DENG, J. On optimizing the backoff interval for random access schemes. *Communications, IEEE Transactions on* 51, 12 (2004), 2081–2090.
- [13] HE, Y., YUAN, R., SUN, J., AND GONG, W. Semi-Random Backoff: Towards resource reservation for channel access in wireless LANs. *IEEE ICNP (2009)*.
- [14] JACKSON, C. Dynamic sharing of radio spectrum: A brief history. In *IEEE DySPAN (2005)*.
- [15] JARDOSH, A., MITTAL, K., RAMACHANDRAN, K., BELDING, E., AND ALMEROOTH, K. IQU: practical queue-based user association management for WLANs. In *MOBICOM (2006)*.
- [16] JARDOSH, A., RAMACHANDRAN, K., ALMEROOTH, K., AND BELDING-ROYER, E. Understanding congestion in IEEE 802.11 b wireless networks. In *ACM IMC (2005)*.
- [17] LIM, C., AND CHOI, C. TDM-based coordination function (TCF) in WLAN for high throughput. In *GLOBECOM (2005)*.
- [18] MANWEILER, J., SANTHAPURI, N., SEN, S., ROY CHOUDHURY, R., NELAKUDITI, S., AND MUNAGALA, K. Order matters: transmission reordering in wireless networks. In *Proceedings of the 15th annual international conference on Mobile computing and networking (2009)*, ACM, pp. 61–72.
- [19] NANDAGOPAL, T., KIM, T., GAO, X., AND BHARGHAVAN, V. Achieving MAC layer fairness in wireless packet networks. In *ACM MOBICOM (2000)*, pp. 87–98.
- [20] PANG, Q., LIEW, S., LEE, J., AND LEUNG, V. Performance evaluation of an adaptive backoff scheme for WLAN. *Wireless Communications and Mobile Computing* 4, 8 (2004), 867–879.
- [21] RAHUL, H., EDALAT, F., KATABI, D., AND SODINI, C. Frequency-aware rate adaptation and MAC protocols. In *MOBICOM (2009)*.
- [22] RHEE, I., WARRIER, A., AIA, M., AND MIN, J. ZMAC: A Hybrid MAC for Wireless Sensor Networks. *ACM SenSys (2005)*.
- [23] ROMAN, B., STAJANO, F., WASSELL, I., AND COTTINGHAM, D. Multi-carrier burst contention (MCBC): Scalable medium access control for wireless networks. In *IEEE WCNC (2008)*.
- [24] SEN, S., CHOUDHURY, R., AND NELAKUDITI, S. Listen (on the Frequency Domain) Before You Talk. *ACM HOTNETS (2010)*.
- [25] SEN, S., ROY CHOUDHURY, R., AND NELAKUDITI, S. CSMA/CN: Carrier Sense Multiple Access with Collision Notification. *ACM MOBICOM (2010)*.
- [26] SEN, S., SANTHAPURI, N., CHOUDHURY, R., AND NELAKUDITI, S. AccuRate: Constellation Based Rate Estimation in Wireless Networks. In *NSDI (2010)*.
- [27] SHIN, S., AND SCHULZRINNE, H. Experimental measurement of the capacity for voip traffic in ieee 802.11 wlans. In *INFOCOM (2007)*, IEEE.
- [28] SHRIVASTAVA, V., AHMED, N., RAYANCHU, S., BANERJEE, S., KESHAV, S., PAPAGIANNAKI, K., AND MISHRA, A. Centaur: Realizing the full potential of centralized wlans through a hybrid data path. In *MOBICOM (2009)*.
- [29] TAN, K., FANG, J., ZHANG, Y., CHEN, S., SHI, L., ZHANG, J., AND ZHANG, Y. Fine Grained Channel Access in Wireless LAN. *ACM SIGCOMM (2010)*.
- [30] VERKAIK, P., AGARWAL, Y., GUPTA, R., AND SNOEREN, A. Softspeak: making VoIP play well in existing 802.11 deployments. In *NSDI (2009)*.
- [31] VUTUKURU, M., BALAKRISHNAN, H., AND JAMIESON, K. Cross-layer wireless bit rate adaptation. *ACM SIGCOMM (2009)*.