

Interference Cancellation: Better Receivers for a New Wireless MAC

Daniel Halperin* Josephine Ammer† Thomas Anderson* David Wetherall‡

Abstract

We argue that carrier sense in 802.11 and other wireless protocols leads to scheduling decisions that are overly pessimistic and hence waste capacity. As an alternative, we propose *interference cancellation*, in which simultaneous signals are modeled and decoded together rather than treating all but one as random noise. This method greatly expands the conditions under which overlapping transmissions can be successfully received, even by a single receiver. We demonstrate the practicality of these better receivers via a proof-of-concept experiment with USRP software radios. We argue that supporting concurrent transmissions enables new and more effective wireless MACs in which carrier sense is disabled.

1 Introduction

Carrier sense, the deferral of transmission while a device senses another in progress, is the dominant mechanism used in wireless LANs (802.11) to share the spectrum between bursty traffic sources. Its use derives from the signal-to-interference-plus-noise ratio (SINR) model of communications theory, which states that wireless data can be successfully received when the signal power (S) sufficiently exceeds the combined power of all interfering transmissions (I) and noise (N). By eliminating interfering transmissions, data will be received over as large a range as possible, with higher SINR permitting the reception of higher rate transmissions. Thus sources can send on demand and as rapidly as feasible.

But carrier sense is inherently wasteful because it discourages spatial reuse. Receivers typically have a large dynamic range, upwards of 60dB in most production systems, and so a single transmission can be detected over a large area. This causes a widespread communication blackout whenever anyone is sending, the equivalent of allowing only a single person to speak at a time in a crowded room. Instead, it is well known that multiple simultaneous conversations can occur safely in practice provided that the right transmitter’s signal is strong at the right receiver, e.g. [14, 17]. That is, carrier sense can be overly pessimistic.

There have been various attempts to address this problem by tuning carrier sense to discover opportunities for

spatial reuse and hence obtain greater capacity. Many techniques raise the carrier sense threshold so that a significantly attenuated transmitter judged to be “far enough away” from a potential sender will not needlessly prevent the node’s transmission [26, 29]. But this approach misses the mark and is fragile because carrier sense is based on conditions at the sender while successful transmission depends on conditions at the receiver. To maximize spatial reuse, recent work tunes not only the carrier sense threshold but also transmit power and data rate [14, 26]. Tuning schemes to date often depend on accurate location information [19], specific node topologies [17, 22], distributed coordination [14, 19, 26, 22], or other factors that are problematic for deployment. Complementary techniques that probe conditions at the receiver such as RTS/CTS [1, 5, 12] have not proven effective in practice either. For example, RTS/CTS is typically disabled [2, 25] because its costs outweigh its benefits in most scenarios.

In this paper, we propose *interference cancellation* [21] as a radical alternative to carrier sense. It enables receivers to function well when signals are transmitted at the same time by different senders. This reduces the need for carrier sense or other forms of scheduling. The key insight we use is that conventional 802.11 receivers treat interfering signals as *random noise*, when in fact they are not noise but *highly structured* signals that correspond to modulated data. By modeling the structure of an interfering signal and then subtracting it out, we greatly reduce the negative impact of interference. At a high-level, this idea is similar to multi-user systems such as CDMA [24], but at a low-level it differs significantly because the system is not designed with codes, synchronization and power control to facilitate the separation of concurrent signals. The technique increases our ability to recover *all* simultaneous interfering transmissions, even at a single receiver and when the SINR before cancellation renders it impractical to receive *any* of the signals with traditional methods.

In wireless networks based on these techniques, we believe that receivers will be able to decode concurrent transmissions often enough to alter the common case MAC to forego carrier sense altogether, especially given that contention levels are typically low in most settings [15]. The result should be a substantial increase in capacity because pessimistic scheduling decisions have been eliminated. Of course, we will not always be able to successfully decode multiple, overlapping transmissions.

*Dept. of Computer Science and Engineering, Univ. of Washington

†Dept. of Electrical Engineering, Univ. of Washington

‡Intel Research Seattle

But with mechanisms such as exponential backoff to handle exceptions, we will no longer be dependent on carrier sense. The limiting factor of our technique is no longer SINR, but k -SNR, the number of interfering signals k that can be received or factored out at a given noise level.

The rest of this paper is organized as follows. In Section 2, we explain a basic technique for interference cancellation and demonstrate its feasibility using USRP software radios. Then, in Section 3, we discuss the implications of interference cancellation for MAC design in 802.11 and emerging wireless systems. We discuss research challenges for deployment in Section 4 and place our work in the context of the large corpus of existing research in Section 5. Finally, we conclude in Section 6 with a discussion of the future research agenda that we propose.

2 Interference Cancellation

In this section we explain a method to model and cancel interference in simultaneous wireless transmissions.

2.1 Theoretical Background

Consider the signal R received containing two overlapping transmissions R_1 and R_2 which, due to the superposition principle, is simply the sum of the two transmissions, plus noise at the receiver:

$$R(t) = R_1(t) + R_2(t) + n(t). \quad (1)$$

A received transmission R_i , converted to baseband from the center frequency f_c by analog frontend hardware, can be represented as

$$R_i(t) = H * \left[h * A(t) \cos[2\pi t(\gamma + f(t)) + \phi(t)] \right] \quad [* \text{ is convolution}] \quad (2)$$

where $A(t)$, $f(t)$, and $\phi(t)$ are the respective amplitude, frequency, and phase shifts determined by the data being sent and the modulation scheme. Due to nonideal physical hardware oscillators that do not perfectly replicate f_c at both transmitter and receiver, conversion to baseband at the receiver will introduce a small frequency offset γ . Also, transmitters and receivers use matched pulse shaping filters h which distort the signal slightly to control spectral use for FCC regulations and to overcome intersymbol interference. Finally, H is the transfer function of the channel which includes effects such as attenuation, multipath, and Doppler shift.

Another common signal representation uses Euler's formulas to rewrite a signal $A(t) \cos(2\pi(\gamma + f(t))t + \phi(t))$ after conversion to baseband as

$$A(t) \cos(\Phi(t)) - jA(t) \sin(\Phi(t)), \quad (3)$$

where $\Phi(t) = 2\pi(\gamma + f(t))t + \phi(t)$ is the instantaneous phase of the baseband signal. We define I to be the real part of the signal (*In phase* with the original wave) and Q to be the imaginary part (at a $\frac{\pi}{2}$ offset, or *Quadrature*).

We use this form to define modulations in terms of a ‘‘constellation’’ mapping bits into waveforms. For example, binary phase shift keying (BPSK) is a simple modulation scheme that maps a single bit b into a phase shift of $b\pi$, sending $\langle 1 + 0j \rangle$ for $b = 0$ or $\langle -1 + 0j \rangle$ when $b = 1$. Symbol decoding is then in principle a simple task: using maximum likelihood detection, a receiver will decode a received sample to the closest point in the constellation after compensating for γ .

2.2 Interference Cancellation

We assume a communication system where f_c , the modulation scheme, and the transmitter's pulse shaping filter h are known. Additionally, we assume for simplicity that the channel H induces only attenuation and a phase shift related to the distance between transmitter and receiver, and that both are constant on the time-scale of a packet. We can exploit periods of isolated transmissions to estimate the average signal amplitude \bar{A} and γ using standard synchronization techniques [16], and more advanced techniques can determine these parameters even under interference.¹ With all of this information and a received signal R_i , we can recreate the original signal S_i .

The estimates of these parameters for both R_1 (with an m -symbol constellation) and R_2 (with an n -symbol constellation) will yield a combined mn -symbol constellation that consists of all pairwise sums of points in the constellations of R_1 and R_2 (Figure 1(c)). This joint constellation varies with time and the frequency offsets of each signal. We can decode both signals at once by mapping a received sample to a *pair of symbols* corresponding to the closest of the mn joint constellation points. This method will be effective as long as the space between the combined constellation points is greater than the noise level. Note that *nothing fundamentally limits us to decoding only two simultaneous signals*, and that for large mn [op . . .], we may not need to explore the entire space.

Figure 1 shows a simulated example of this technique using BPSK. Figure 1 (a) shows the amplitude of received data where at first only S_1 transmits, S_2 begins to send at $t = 500$ and both signals are received, and then at $t = 1000$ S_1 stops sending and S_2 transmits alone. Plots (b), (c), and (d) each display a scatterplot of the values that might be sampled at a specific time during the transmission, in the complex plane. (The relative orientation of the two pairs will vary with time due to different values of γ .) Parts (b) and (d) show symbols received under noise in the constellations of S_1 and S_2 , and in (c) we see the joint constellation consisting of four clusters for the four possible values sent by S_1 and S_2 . In the case shown, a simple detector may work most of the time, even in the presence of interference; by modeling the two signals explicitly, we can reduce the bit error rate to result in a practical system.

¹For example, see Katti et al. [13] for a method to estimate \bar{A} .




Figure 1: Synthetic illustration of interference cancellation. The (a) amplitude of the received data enables us to differentiate the received samples as S_1 sending for $t \in [0, 1000]$ and S_2 during $[500, 1500]$. Scatterplots (b), (c), and (d) show noisy constellations in I/Q format at one particular time instant each, during periods without interference (b and d) and combined (c). (In general, the relative angle of the constellations will vary.) A simple detector aware of only one signal may decode many bits correctly if locked onto the strong signal, but modeling the two signals explicitly allows us to decode the four clusters in (c) correctly for both signals.

2.3 Proof of Concept

We have demonstrated the practicality of interference cancellation using the Universal Software Radio Peripheral [7] with the GNU Radio libraries [8]. In a three-node USRP testbed, two nodes S_1 and S_2 sent DBPSK-encoded random frames at a low data rate (125kbps) with a short random interval between sends, while a third receiver node sampling at $8\times$ the symbol rate (1MHz) logged the raw measured electromagnetic waves. We then replayed the stored signals with a detector using standard techniques² and a detector employing interference cancellation. This allows for a fair, reproducible comparison between the two detection methods.

Each node sent 1000 packets with 256-byte payloads including a 32-bit CRC. The stronger S_1 had an SNR of 29.8 dB and the SNR of S_2 is 25.9. In isolation, both senders will receive packets with high probability (respective packet error rates of $1.17e-10$ and $5.8e-9$). However, when both senders transmit the SINR of S_1 is 3.4dB predicting a BER of 0.017 using the expected error performance for DBPSK [20]. Thus the transmissions of S_1 will not succeed (probability of success is $4.2e-15$ if the packets completely overlap) when S_2 interferes.

Table 1 shows the results of our simple experiment. Of 2000 packets transmitted, 398 experienced no interference and were received by both detectors. The simple detector synchronized on (successfully decoded the header of) the first packet in 91.5% (733/801) of collisions but the joint detector locked onto 86.3% (1383/1602) of *both* packets involved in collision, illustrating that with joint detection, synchronization can be achieved even under interference.

Note that none of S_2 's interfered-with packets passed the CRC check, and only 42 (5.2%) of S_1 's did. Interference cancellation received 137 (17.1%) of S_2 's packets and 483 (60.3%) of S_1 's under interference. This alone is

²As the default GNU Radio DBPSK detector has poor error performance, we re-implemented it for a more meaningful comparison.

	Traditional Receiver	Interference Cancellation
Total packets:	2000	
Isolated transmissions:	398	
Collisions:	801	
Of 1602 Interfered-With Packets		
Packets synchronized:	733	1383
S_1 :	369	731
S_2 :	364	652
BER (sync packets):	0.0897	0.0200
S_1 :	0.0038	0.0070
S_2 :	0.1768	0.0346
Correct bits (in pkts):	667.2	1355.3
S_1 :	367.6	725.9
S_2 :	299.6	629.4
Total CRC pass:	42	620
S_1 :	42	483
S_2 :	0	137
CRC pass 0.5% FEC:	266	1139
S_1 :	266	694
S_2 :	0	445

Table 1: Table of reception rates

not enough to improve performance over a single sender, but even 0.5% FEC (10 errors in 2048 bits!) dramatically improves recovery such that 71% (1139/1602) of the colliding packets are recovered. This is an increase of 42% over the single sender case, and suffices to demonstrate at least one realistic situation in which interference cancellation improves throughput.

2.4 Limitations

When would interference cancellation work? There are theoretical and practical limits that we address briefly here, leaving full quantification as future work.

To begin, consider two equal-amplitude, perfectly synchronized senders using the same modulation. When these senders transmit different symbols, the receiver has no means to distinguish which node sent which symbol – either possibility yields the same sample. Generally, there

can be many situations in which two points in the combined constellation can be close together, leading to ambiguity in some of the decoded bits. The frequency offset γ causes the constellations to rotate relative to one another and prevents such synchronization from happening except through periodic alignment. The density of the combined constellation and the noise level determine how much ambiguity exists, and FEC can be used effectively to combat small amounts of unpreventable loss.

Practically, the sampled digitized signal processed by the receiver must contain sufficient information about all interfering signals to recover their bits. Precision can be lost in the digitization process as the analog signal is truncated; 10-bit digital samples cannot represent a signal that is a factor of 1024 weaker than the strongest interferer. However, modern analog-to-digital converters have sufficient precision to represent a large dynamic range. In addition, the automatic gain control (AGC) loop which normalizes signal amplitude before digitization must amplify an incoming signal to the right level such that both strong and weak interfering signals can be recognized and decoded.

3 Implications for MAC Design

We have introduced interference cancellation and demonstrated its feasibility with a proof of concept implementation. Note that using interference cancellation increases the total number of bits received at the cost of a higher BER compared to isolated transmissions, requiring a higher error correction rate to be robust to multiple receivers. An important question to consider is *Have we just exchanged a single fast link for multiple slow ones?*

Wireless radios in 802.11 operate many dB above the theoretical limits for information transfer – a limited set of data rates prevent the Shannon capacity from being achieved. Software radios and improved engineering techniques may lead to systems with dynamic bitrates able to operate closer to link capacity. However, we believe that even in such a future there will be a non-trivial link margin to permit packet synchronization, and for robustness to noise and fading (Intersil recommends an extra 30dB over the minimum SNR for 10^{-6} BER [30]). With multiple sender-receiver pairs, in a mesh or with multiple nearby APs, a transmission is unlikely to be exactly at the limit for an unintended receiver – and, the scenario where a receiver desires to receive only one of the packets requires a lower margin. These factors combine to imply that there is additional capacity which other devices can utilize to send concurrently – so we may exchange a single high capacity link for multiple slower links that combined carry more net information.

A second benefit of disabling carrier sense is that a major limiting factor in wireless performance is MAC overhead in terms of poor channel utilization – in 802.11,

with one sender, the medium has been observed to be idle 70% of the time [15]. Even with multiple senders, using 1500-byte payloads (most wireless links will never exceed the TCP MTU), the maximum throughput of 802.11b “11Mbps” is actually 6.1Mbps and 802.11a’s “54Mbps” rate reaches about 32Mbps [11]. Especially since contention is typically low [15], a MAC that enables more air-time utilization in the common case could lead to a huge performance boost.

That brings us to the question of *What MAC should be used in a carrier sense free system?* A simple answer is: in the common case, none – let nodes send traffic as desired and use ACK-based binary exponential backoff for congestion control. For graceful recovery when interference cancellation fails, we supplement the basic backoff with mechanisms for a receiver to initiate backoff as well.

What are the concerns with this approach? One obvious concern is that a loud transmitter will drown out weaker senders due to limited precision as discussed in Section 2. To combat this, suppose that the receiver can detect a weak ongoing transmission, even if unable to decode that packet. Then it can delay ACKs to the strong transmitter or include in its ACK a “Request-to-Receive” (RTR) asking it to back off. Note that, in a way, this approximates a sort of “receiver sense.” When the receiver cannot even detect the presence of this secondary signal, then there’s not a clear solution other than embedding artificial backoff in the default protocol. On the other hand, if the loud transmitter is so much stronger than the weaker, perhaps the right solution is for it to send all its data at its fast data rate and then stop, compared to 802.11’s packet-level sharing. The right definition of fairness in such imbalanced wireless congestion situations is a yet-unresolved question, but the ability of interference cancellation to support simultaneous slow and fast senders can help solve the slow sender bottleneck observed for example in networks with coexisting 802.11b and g devices [6].

The opposite case is many weak transmitters, all just at the limits of the receiver’s ability to successfully decode, interfering with each other. Again, evidence of low contention [15] makes this scenario unlikely, but in addition this is exactly the case for which exponential backoff is designed and should resolve to an efficient steady state.

The next question that follows is *How should this technology be deployed?* Unlike most other work [5, 13, 14, 19, 22, 27] in this area that requires mesh networks, complex protocol modifications, and universal participation to realize its benefits, implementing interference cancellation techniques unilaterally at a single receiver will make it more robust. One key is that this technology fits directly into the access point usage model (as well as in the mesh network model) – only the hardware has to change. To switch the network over to carrier sense free operation, this technology can be deployed first at access points

and then in wireless NICs. An access point, as the central location through which all traffic must pass, will be able to recognize the presence of legacy clients and using delayed ACKs and RTR packets can force the next-generation clients into leaving room for legacy clients to communicate.

An important concern is *how interference cancellation will interact with more advanced technologies* as they become commercialized. The next generation of wireless will be marked by MIMO [3], in which multiple antennas are used at both sender and receiver to transfer different simultaneous streams of data along the separate paths between pairs of antennas. However, the superposition principle applies to MIMO as well, and interference cancellation techniques are orthogonal to MIMO and will still apply when it is deployed.

An important part of the wireless future will be the coexistence of heterogeneous devices. Today Bluetooth, 802.11, 802.15.4 (e.g. Zigbee), RFID, and other devices such as cordless phones and microwaves all operate in the same part of the spectrum, but in practice there are adverse interactions in their interoperation [9]. We believe that the future of wireless will be marked by further heterogeneity in the ISM bands, and the systematic cancellation of interference will become critical for good performance in the presence of heterogeneous interfering devices.

4 Challenges for Interference Cancellation

What research challenges do we face in realizing the potential of interference cancellation?

First, we need to understand when interference cancellation works and when it will not. In Section 2 we presented a brief discussion of a few challenges, but in order to characterize the limits of its functionality in terms of complex modulations, required link margin, and robustness to multiple senders, comprehensive experimental evaluation will be required. Part of this work will include adapting signal processing algorithms (such as AGC and packet synchronization) that assume a single sender.

Secondly, how should the physical layer be defined to maximize recovery? For instance, some modulation schemes are better suited to concurrent transmissions than others, and it's not immediately clear which are optimally suited for our MAC. Rectangular QAM-16 and 16-PSK both carry 4 bits per symbol; 16-PSK has worse error performance in the single sender case [20], but its joint constellation is less ambiguous. Other physical layer changes might include randomization on a sender's part to desynchronize its symbols from an interferer's.

Next we need to consider interactions with FEC and other error recovery techniques. Spreading and error correction will enhance joint detection by correcting wrong symbol decodings and lowering ambiguity by reducing the possible symbol combinations. As well, confidence

information from maximum likelihood detection can be passed to higher-level error recovery systems such as PPR [10] that take advantage of it.

Finally, how should the MAC work in detail? Our suggestions provide low overhead and maximum utilization in the common single-sender case, but the subtleties of when interference cancellation works will be important in designing a MAC to gracefully handle contentious environments. Additionally, understanding the right receiver methods for characterizing the channel, and for using channel feedback, must be researched.

5 Related Work

There is a large body of related work in the area of wireless communications. Many existing reception techniques [28] use antenna arrays to enable the reception of multiple concurrent transmissions. In contrast, single antenna interference cancellation (SAIC) techniques like that described here work without an array and as such are deployable in a strict superset of the environments for which multi-array reception techniques are designed.

Interference cancellation [21] is an example of blind signal separation (BSS) [4], one of many techniques used for multi-user detection (MUD) [23]. At a high level, our proposal for utilizing it in wireless LANs is like this existing work, but differs substantially at a lower level. MUD work in CDMA [24] networks exploits dedicated frequency bands and specific structure added to the overall system, e.g., different CDMA codes for different users, cells that separate areas of reception, and centralized node synchronization and power control. Work deploying SAIC [18] has similarly focused on deployment on cell phones in GSM networks to enable them to isolate a transmission from a single cell tower. In contrast, nodes in our systems may all use the "same code," have varying frequency offsets with respect to one another, are uncoordinated and may be in networks that completely overlap in their coverage range. Operating in the ISM band, there is no central administrator controlling which devices use the band or that provides isolation and power control.

The details of interference cancellation are similar to but more general than recent work on analog network coding in which two signals are combined in a transmission and later separated [13, 27]. An important difference is that the information in one of the signals is known and then used to decode the other signal. In our case, neither signal is known before decoding. Katti et al. [13] suggest that it would be possible to decode the stronger signal, if its SINR is sufficient, and then remove that signal and decode the weaker signal, but this proposal is less effective than the joint signal detection we use to receive even when standard techniques do not recover the stronger signal.

Finally, there is work on improving carrier sense. Mechanisms like RTS/CTS [12] help in some corner cases

but have proven to hurt performance in practice [25]. Clear channel assessment adaptation methods [14, 29] help achieve better spatial reuse but are fragile in real systems. Moscibroda et al. [17] show through careful node placement and power control that sending packets without carrier sense can greatly increase channel capacity. They observed throughput gains of nearly $2\times$ the theoretical single-transmitter limit and nearly $3\times$ over default CSMA. Without interference cancellation to reduce the effects of interference, this work relied on precise location knowledge and careful topology planning, and was dramatically sensitive to location changes of even a few centimeters. In contrast, our proposal for leveraging interference cancellation in wireless LANs is more robust than this related work, does not require modifications to the existing WLAN models, has a clear picture for deployment, and is designed to work in the chaotic ISM band.

6 Conclusions

This paper proposes redesigning the MAC for 802.11-like WLANs to support concurrent transmissions via interference cancellation as a radical alternative to serializing communication with carrier sense. The potential benefits of this shift are twofold with increased capacity through more aggressive spatial reuse and more efficient airtime utilization in the common, low contention cases. Looking ahead, as the population and density of coexisting and competing heterogeneous devices in the wireless ISM bands grows, interference cancellation becomes increasingly important in maintaining system operation.

We demonstrate through a simple proof of concept experiment that interference cancellation is feasible in practice but much work stands between what we have today and the realization of an interference robust MAC. Work both theoretical and practical is required as signal processing algorithms need to be adjusted to this new model and hardware designed to support it. After the groundwork has been laid, the main challenge in this system is to gain a better understanding of where interference cancellation works, where it will not, and the gains it can provide. Only through the deployment and evaluation of this technology can it be fully understood.

Acknowledgments

We would like to thank the anonymous reviewers for their comments. This work was supported in part by the National Science Foundation (CNS-0435065), Intel, Microsoft, and the UW College of Engineering.

References

[1] V. Barghavan, A. Demers, S. Shenker, and L. Zhang. MACAW: A media access protocol for wireless LAN's. In *ACM SIGCOMM*, 1994.

[2] J. Bicket, D. Aguayo, S. Biswas, and R. Morris. Architecture and evaluation of an unplanned 802.11b mesh network. In *ACM MobiCom*, 2005.

[3] E. Biglieri, R. Calderbank, A. Constantinides, A. Goldsmith, A. Paulraj, and H. V. Poor. *MIMO Wireless Communications*. Cambridge University Press, 2007.

[4] J.-F. Cardoso. Blind signal separation: statistical principles. *Proceedings of the IEEE*, 9(10):2009–2025, 1998.

[5] M. Cesana, D. Maniezzo, P. Bergamo, and M. Gerla. Interference aware (IA) MAC: an enhancement to IEEE 802.11b DCF. In *IEEE VTC*, 2003.

[6] Y.-C. Cheng, J. Bellardo, P. Benko, A. C. Snoeren, G. M. Voelker, and S. Savage. Jigsaw: Solving the puzzle of enterprise 802.11 analysis. In *ACM SIGCOMM*, 2006.

[7] Ettus Research, LLC. <http://www.ettus.com>.

[8] GNU Radio Project. <http://gnuradio.org/trac>.

[9] I. Howitt, V. Mitter, and J. Gutierrez. Empirical study for IEEE 802.11 and bluetooth interoperability. In *IEEE VTC*, 2001.

[10] K. Jamieson and H. Balakrishnan. PPR: Partial packet recovery for wireless networks. In *ACM SIGCOMM*, 2007.

[11] J. Jun, P. Peddabachagari, and M. Sictiu. Theoretical maximum throughput of IEEE 802.11 and its applications. In *IEEE Net. Comput. and Apps. (NCA)*, 2003.

[12] P. Karn. MACA - a new channel access method for packet radio. In *ARRL 9th Computer Networking Conf.*, 1990.

[13] S. Katti, S. Gollakota, and D. Katabi. Embracing wireless interference: Analog network coding. In *ACM SIGCOMM*, 2007.

[14] T.-S. Kim, H. Lim, and J. C. Hou. Improving spatial reuse through tuning transmit power, carrier sense threshold, and data rate in multihop wireless networks. In *ACM MobiCom*, 2006.

[15] R. Mahajan, M. Rodrig, D. Wetherall, and J. Zahorjan. Analyzing the MAC-level behavior of wireless networks in the wild. In *ACM SIGCOMM*, 2006.

[16] H. Meyr, M. Moeneclaey, and S. A. Fechtel. *Digital Communications Receivers: Synchronization, Channel Estimation, and Signal Processing*. John Wiley & Sons, 1998.

[17] T. Moscibroda, R. Wattenhofer, and Y. Weber. Protocol design beyond graph-based models. In *HotNets-V*, 2006.

[18] A. Mostafa, R. Kobylinski, I. Kostanic, and M. Austin. Single antenna interference cancellation (SAIC) for GSM networks. In *IEEE VTC*, 2003.

[19] T. Nadeem, L. Ji, A. Agrawala, and J. Agre. Location enhancement to IEEE 802.11 DCF. In *IEEE INFOCOM*, 2004.

[20] J. G. Proakis. *Digital Communications*. McGraw-Hill, 2001.

[21] H. Schoeneich and P. A. Hoeher. Single antenna interference cancellation: Iterative semi-blind algorithm and performance bound for joint maximum-likelihood interference cancellation. In *GLOBECOM*, 2003.

[22] A. Vasani, R. Ramjee, and T. Woo. ECHOS - enhanced capacity 802.11 hotspots. In *IEEE INFOCOM*, 2005.

[23] S. Verdú. *Multuser Detection*. Cambridge University Press, 1998.

[24] A. J. Viterbi. *CDMA: Principles of Spread Spectrum Communication*. Addison-Wesley, 1995.

[25] K. Xu, M. Gerla, and S. Bae. Effectiveness of RTS/CTS handshake in IEEE 802.11 based ad hoc networks. In *Ad Hoc Network Journal*, 2003.

[26] X. Yang and N. Vaidya. On physical carrier sensing in wireless ad hoc networks. In *IEEE INFOCOM*, 2005.

[27] S. Zhang, S.-C. Liew, and P. P. Lam. Physical-layer network coding. In *ACM MobiCom*, 2006.

[28] P. X. Zheng, Y. J. Zhang, and S. C. Liew. Multipacket reception in wireless local area networks. In *IEEE ICC*, 2006.

[29] J. Zhu, X. Guo, L. L. Yang, W. S. Conner, S. Roy, and M. M. Hazra. Adapting physical carrier sensing to maximize spatial reuse in 802.11 mesh networks. *Wirel. Comm. Mob. Comput.*, 4(8):933–946, December 2004.

[30] J. Zyren and A. Petrick. Tutorial on basic link budget analysis. Intersil Corp., <http://www.sss-mag.com/pdf/an9804.pdf>, 1998.