

Design Considerations for Minimal-Power Wireless Spread Spectrum Circuits and Systems

BRENT A. MYERS, MEMBER, IEEE, J. BARTOW WILLINGHAM, MEMBER, IEEE, PATRICK LANDY, MARK A. WEBSTER, MEMBER, IEEE, PERRY FROGGE, AND MICHAEL FISCHER, MEMBER, IEEE

Wireless local area networks (WLANs) provide untethered access to digital data in both home and office environments. To preserve mobility for extended periods of time, low power design is a critical enabler. Low-power WLAN design requires consideration at every level from the system topology, protocol, and radio architecture to the detailed circuit design. Every level is required to work together in harmony to gain the maximum in power reduction. This paper discusses the performance requirements and design tradeoffs encountered in the design of minimal-power WLAN systems.

Keywords—Local area networks (LANs), low-power IC design, RF IC design, wireless communications, wireless data communications, wireless IC design, wireless local area networks (WLANs), wireless networking.

I. INTRODUCTION

In this paper, design considerations for wireless local area networks (WLANs) will be considered, specifically those networks designed to IEEE Standard 802.11, which has emerged as a leading standard for high data rate applications for both office and home environments. Section II will briefly overview the concepts behind wireless networking as well as the evolution of the 802.11 standard. Section II will also overview the architectural and partitioning considerations which influence low-power operation, and discuss high level concerns in the three main components of the system including the front-end radio, digital baseband processor (BBP), and the media access control (MAC) unit, which interfaces with the host. Sections III–V will then discuss architectural and implementation details of each of the main components as applied to a true system for up to 11 Mb/s capability. Section VI completes the paper with a look into the future.

II. BACKGROUND

The concept of a local area network allows the use of wired or wireless media. However, the layering of LAN

protocols defined in IEEE Standard 802-1990 [1] makes several assumptions about the characteristics of the physical layer, which are valid for all wired media but not for a wireless physical layer. These assumptions include full connectivity among the stations connected to a given LAN, physical isolation between co-located, independent LANs, the existence of a readily discernable network boundary that is under the control of the organization operating the LAN, and a static network topology during the exchange of protocol data units (PDUs) over the LAN. The protocols used at and above the logical link control (LLC) sublayer for peer-to-peer communication over a LAN implicitly assume that the active stations are continuously accessible, that utility power is available at the physical locations served by the LAN, and that the bit error rate (BER) for communication over the LAN is better than 10^{-6} . The MAC sublayer of IEEE Standard 802.11 has to render these physical layer differences transparent when viewed from the MAC service access point. In order to provide access control for the shared medium, addressing and recognition of frames in support of LLC, frame check sequence generation and checking, and LLC PDU delimiting, as required of any IEEE 802 MAC sublayer, a number of additional functions are needed to compensate for the different physical layer characteristics. This results in IEEE 802.11 having one of the most complex MAC protocols ever standardized. While the assumption of utility power is the focus of this paper, all of the other assumptions and the techniques to make them valid to the higher layers come into play as additional functions requiring power and as such threatening to shorten battery life.

Key system concepts that impact power are the cell size, the physical layer transmission rate, the overhead associated with the transmission (the efficiency of the protocol), the ability to put the receiver to sleep, speed of wake-up, the efficiency of the modulated waveform (as measured by E_b/N_o at a particular BER), antenna patterns, acquisition aids, system frequency uncertainty and methods of combating multipath, to name a few. In a minimum power system design we strive to transmit as little as possible with low transmit power and to have a low-power receiver that we can turn off completely periodically.

Manuscript received December 31, 1999; revised June 21, 2000.

The authors are with the Intersil Corporation, Melbourne, FL 32902-0883 USA (e-mail: bam@eng.intersil.com; bwilli08@intersil.com; pjl@eng.intersil.com; mwebster@intersil.com; pwf@eng.intersil; mfischer@choicemicro.com).

Publisher Item Identifier S 0018-9219(00)09134-9.

0018-9219/00\$10.00 © 2000 IEEE

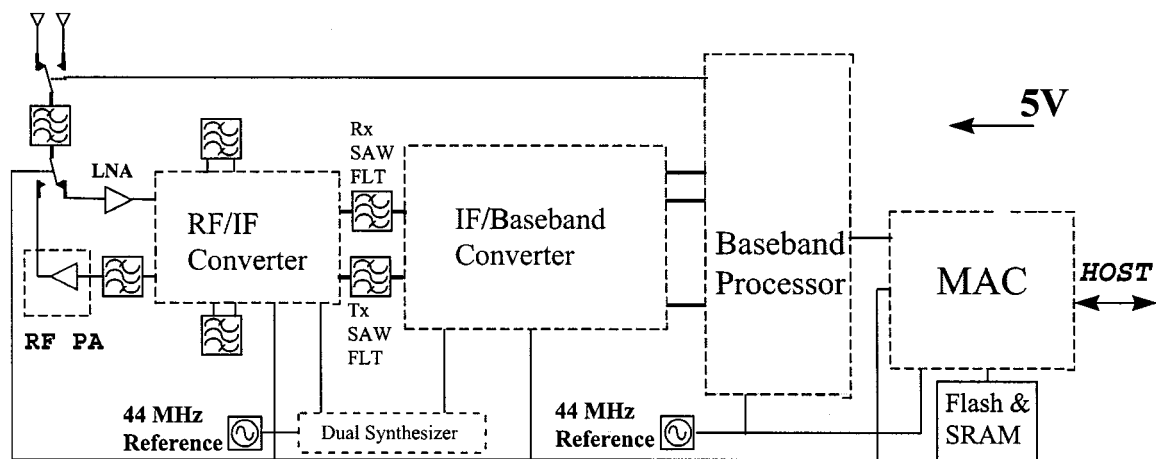


Fig. 1. Prism I radio architecture for up to 2 Mb/s data rate.

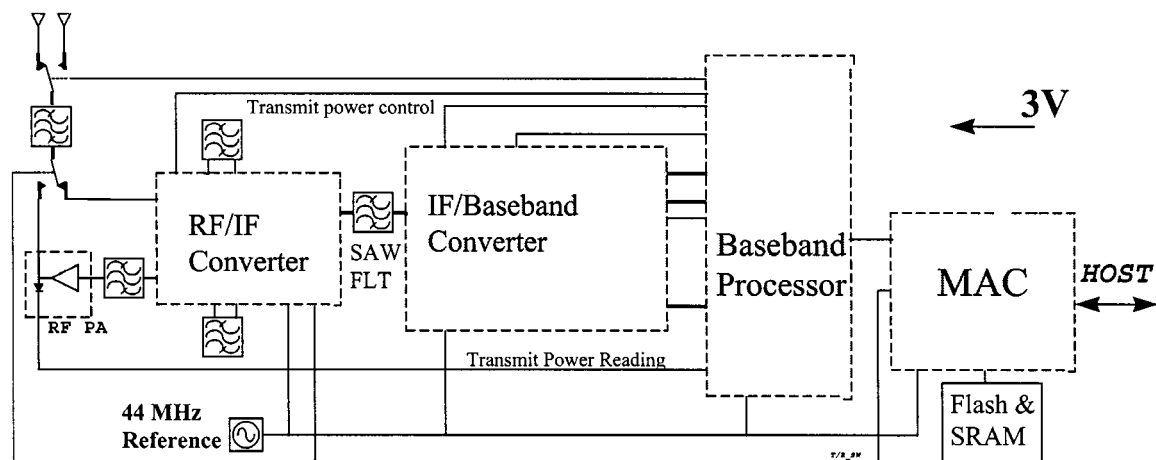


Fig. 2. Prism II radio architecture to support up to 11 Mb/s (802.11 b).

The IEEE 802.11 standard has been constructed to support low-power operation. IEEE 802.11 is targeted at cell sizes on the order of hundreds of feet with transmission rates of 1, 2, 5.5, and 11 Mb/s using efficient phase shift keying (PSK) waveforms with processing gain and multipath protection. The standard allows packet mode transmission at lengths up to 2304 bytes with headers to aid in the acquisition of the carrier and symbol timing. The use of PSK waveforms together with the relatively small cell sizes allows power amplifiers to be on the order of only 20 mW. The output power requirement is a significant driver for total system power. In order to reduce power, the goal is to reduce the time that the transmitter is on and to transmit with as low a power as possible consistent with the range and interference levels.

The requirement of 25 ppm oscillators allow acquisition without sweeping frequency or timing, thus adding to the efficiency of the protocol without increasing the cost unnecessarily. This efficiency gain is possible because of the reduced preamble used for acquisition and increases the system capacity for a given physical layer raw data rate.

A. Radio Architectures to Support Low Power

This section describes the radio architectural issues involving power consumption. Two generations of network in-

terface card (NIC) architectures are contrasted to compare the changes made to reduce power. Intersil Corporation (formerly Harris Semiconductor) has developed two network interface cards meeting the requirements of IEEE 802.11. The first is called Prism I and meets the requirements of the original standard published in 1997. The second is called Prism II and implements the high rate extensions developed by the IEEE 802.11 Task Force B. Both cards are supported by an integrated chip set for maximum cost reduction and minimal power. The Prism II radio makes a set of architectural improvements over the Prism I design which reduce power and cost while improving performance. One or both architectures implements the following power and cost savings:

- transmit power control;
- use of a single SAW filter;
- use of a single oscillator;
- half duplex radio with unused portions turned off;
- sleep modes for fast recovery;
- low-power acquisition mode;
- lower voltage operation.

The Prism I architecture is shown in Fig. 1 and the Prism II architecture is shown in Fig. 2. The Prism II network interface card reference design is illustrated in Fig. 3.




Fig. 3. Prism II network interface card.

Transmit power control is implemented in Prism II as an improvement over the previous radio. Transmit power control provides the capability of adjusting the transmitted power at the transmit IF to maintain a constant output power at the antenna port. This increases power efficiency by allowing the power amplifier to operate closer to its compression point without concern for overdriving it and increasing the sidelobes beyond what is allowed in the standard or by the regulatory domains. Without transmit power control, the tolerances in the transmit chain routinely have ± 2 -dB variation over process and another ± 4 -dB variation over temperature. In the absence of transmit power control and in order to be certain that the power amplifier is not overdriven over production quantities, the transmit power must be backed off by at least 4 dB even when the process variation is removed at manufacturing test. To achieve the same power at the antenna, the power amplifier, which is sized for higher transmit power and then backed off by as much as 4 dB, will draw significantly more power than a lower total power amplifier that is operating close to its compression point. While

adding transmit power control adds some additional system complexity, it allows for a significant decrease in transmit power. This is accomplished by providing a power measuring capability at the power amplifier output and feeding back this measurement to a control loop, which steers the transmit power to the correct value.

The use of a single SAW filter for both transmit and receive primarily reduces the cost of the radio with no increase in power.

Prism I radios used two oscillators, one for the RF and IF synthesizers, and a separate oscillator for the baseband processor and MAC. This allows for improved isolation between the RF/IF and digital sections. Through the use of careful layout the Prism II design allows the use of a single oscillator saving both cost and power. The single oscillator provides the reference for the synthesizers, the carrier timing and symbol timing as well as the clock to the MAC. At the system level, this allows the use of a new IEEE 802.11 feature where the baseband processor can take advantage of the fact that the carrier and symbol timing are locked, allowing both to be

tracked. This feature improves performance without sacrificing power.

Since the 802.11 Protocol is half duplex, then only half of the radio circuitry is used at any one time. One obvious savings is to turn off the circuitry which is not being used. The method of accomplishing this is via the media access controller (MAC) which is controlling the radio. Separate power control enable lines are provided for the overall radio, the power amplifier, the receive RF, IF and baseband, and the transmit RF. These control lines allow sections of the circuitry to be put to sleep with rapid recovery when awakened. A series of power states can be defined.

NIC Power Down—All functions are powered down. Recovery time includes starting and initializing the MAC, loading the initial values and starting the radio, synthesizer loading and acquisition, system slot time acquisition, channel scanning, beacon acquisition, system authentication. No power is required.

Radio Deep Sleep—Allows all initialization registers to be remembered, system beacon timing is remembered, system slot timing is remembered, last good channel is remembered, system authentication is remembered. This requires that the MAC be running while the radio is turned off. Only the synthesizers must reacquire on power up.

NIC Sleep Mode—The MAC is clocked with a low rate clock (such as a 32-KHz watch crystal). This allows all initialization registers to be remembered, last good channel is remembered, system authentication is remembered, beacon timing can be maintained. Slot timing is lost, synthesizers must be reloaded and reacquire.

Radio Receive—Synthesizers are never turned off so that switching between transmit and receive is very rapid. All transmit functions except the synthesizers which are used for both transmit and receive are turned off during receive.

Radio Transmit—All receive functions are turned off except the synthesizers which are used for both transmit and receive.

In most IEEE 802.11 networks, the network interface card spends most of its time in receive mode looking for a packet header. Significant power savings can be achieved by:

- a) turning off the receiver periodically;
- b) controlling the acquisition circuitry to reduce power in noise.

IEEE 802.11 accommodates power savings by a defined mode that allows the receiver to turn off between beacons and to be alerted during a beacon that the access point has data for it. This system level power saving mode can be the most effective of all. This is described in more detail in Section V.

By taking advantage of signal detectors in the baseband processor, significant power saving can be achieved by not running the complex portions of the acquisition circuitry except when a likely candidate for a packet header is found. Section IV describes this technique.

Finally, Prism II operates at 3 V while Prism I operates at 5 V. This alone makes a significant power improvement and is discussed further in Section III.

III. FRONT-END PARTITIONING AND CIRCUIT CONSIDERATIONS

Front-end circuitry for the Prism II WLAN system includes everything between the antenna and the data acquisition networks in the digital BBP (Section IV). These are essentially the components that make up the radio portion of the system. Some of the details pertaining to the radio design will be considered in order to understand relevant power tradeoffs.

A. Front-End System Requirements

The simultaneous requirements to support the short preamble capability for high data rate operation and compensate for multipath channel induced distortions create several challenges for the front-end radio design. On the receive side, multipath compensation takes the form of a DS-based RAKE processor in the Prism II architecture. For proper operation, all processing preceding the RAKE must be linear which means the entire RF, IF, and analog baseband sections. Further, due to the large dynamic range of signals that must be handled by the radio, automatic gain control (AGC) circuitry is necessary to maintain linearity throughout the analog signal processing chain. The Short Preamble option severely constrains the allotted time for the AGC system to settle, where $10 \mu\text{s}$ is a typical design limit. The actual AGC circuitry is given an even smaller time to achieve its gain change (under $1 \mu\text{s}$) in order to allow for maximum integration time. The constraint of linear processing affects the receiver design in the very first amplifier stage (a low noise amplifier or LNA) due to the large maximum signal of -4 dbm that the standard requires. The LNA stage is not able to efficiently handle this large signal unless some AGC or attenuation is applied at RF. However, the sensitivity or minimum signal to be processed is -90 dbm , resulting in a large AGC range ($>90 \text{ db}$) that cannot be achieved entirely at RF due to isolation issues. Thus, the requirement for linear processing throughout the receiver prior to the RAKE requires some combination of RF and IF or baseband analog AGC that settles quickly. Gain/AGC partitioning affects the receiver power required because of signal swing and speed of settling, so is a critical concern for low-power design.

Another important factor which will effect radio performance (power) occurs at the analog and digital baseband interface where the baseband analog I/Q signals connect to the I/Q A/D inputs on the receiver side or the I/Q DAC outputs on the transmit side. It is possible to ac couple at this interface, which has the advantage of eliminating dc offsets. However, this will compromise performance and also slow down transmit-to-receive or receive-to-transmit switching times. For these reasons, a dc coupled interface is preferred. Dc coupling requires low offset voltages $<10 \text{ mV}$ to ensure excellent carrier suppression on the transmit side, and is necessary to provide maximum A/D dynamic range on the receive side. These levels of dc offset must be achieved through inherent low offset circuitry or dc correction techniques. Such approaches increase total power fractionally, but are critical for proper operation.




Fig. 4. Prism II system emphasizing the radio architecture.

Analog impairments play a role in any digital communication system, but as the performance of the radio increases, their significance can become critical. The 11-Mb/s system is more sensitive to certain impairments. Quadrature imbalances can cause increased packet error rates (PER). Gain balance of <1 db and phase balance <2 degrees for the entire processing chain are necessary for a high-performance system. Phase jitter on any local oscillator is another sensitive error. System simulations show that integrated phase jitter <3 degrees rms over a 10–750-KHz bandwidth is required for oscillators in the Prism II architecture.

Distortion and spurious responses are critical in all radio designs and are further compounded in these applications by the low supply voltage of 3 V under which all circuits must operate. On the receive side, distortion can be generated by adjacent channel or large out-of-band signals. The linearity of RF circuits as specified by the well-known third-order intercept (IP3) and the selectivity of front-end filters will primarily determine the interference immunity of the radio from these undesired signals. Analog I/Q baseband circuitry have a different set of distortion challenges. In order to cancel large multipath signals, the RAKE processor in the DSP section must linearly process these signals. Requiring linear processing of multiple signals increases the peak to rms (sometimes called crest factor) of the waveform and forces the circuitry to have increased dynamic range, which requires higher power (to assure linearity). Worst case total harmonic distortion requirements are $<1\%$ across the full dynamic range. On the transmit side, the standard specifies a certain spectral mask, which has the familiar filtered $\sin x/x$ characteristic. Any in-band distortions can impact this spectral mask. Power efficient transmit design usually allots the highest power components the greatest distortion, the principal being that power efficiency increases with decreasing linearity. For this reason, the transmit chain should permit the final power amplifier stage to be the least linear. This means all components preceding the power amplifier should add very little distortion. Regulatory bodies like the FCC also limit any spurious emissions, which means

distortion components or other signals must be minimized by a combination of inherent circuit techniques and filtering in the Tx chain.

Noise is a fundamental limitation that is made worse by the 11-Mb/s system due to the increasing E_b/N_o requirement. In order to maintain range, it is desirable considering the link budget that receiver noise figure (NF) be minimized. This low NF must also be achieved in a system that has antenna diversity (two antennas for mobile applications is typical). For a typical receiver chain following the antennas there is a diversity antenna switch, band select filter, and transmit/receive switch before the signal gets to the first LNA stage. These other components preceding the LNA can have typically 3 db of loss resulting in a 3-db NF before the LNA stage. Thus, it can be challenging to achieve a desired system NF, typically <7 db. Further, if enough gain is not generated at RF, the IF stages that follow can contribute significantly to the system NF.

B. Architecture and Circuit Partitioning

Prism II is a single IF conversion architecture. The single IF architecture meets high performance requirements while minimizing external components and power. Alternative architectures like digital IF and zero IF have disadvantages for this application. Digital IF architectures are capable of achieving excellent performance as many analog impairments are eliminated. However, digital IF systems require increased power and complexity due to the fast sampling rates required. Zero IF architectures minimize external components. The disadvantages of zero IF systems are in the areas of performance and practical implementation. The single IF architecture strikes the best balance considering the many tradeoffs and is illustrated in Fig. 4 for Prism II.

Referring to Fig. 4, the front-end receive chain begins at the antennas and ends at the A/D interface in the HFA3861 baseband processor. The transmit chain begins at the D/A interface in the baseband processor and ends at the antennas. The 2.4-GHz ISM band RF signal must be converted first to

an intermediate frequency (IF) and finally to quadrature (I/Q) baseband signals which range from DC-8.5 MHz. The higher the IF frequency selected, the easier it will be for RF filters to attenuate unwanted image and LO signals. However, too high an IF frequency can cause increased power and potential stability issues for high-gain IF stages (e.g., Rx IF AGC). 374 MHz is the IF frequency for Prism II, which resulted from these considerations as well as an analysis of potential mixing products.

Important factors driving front-end circuit partitioning are performance, power, isolation, packaging, cost, and technology. The performance of 2.4-GHz RF circuits are critically influenced by the package selected. For example, the transmit power amplifier stage (HFA3963) delivers about 100 mW of linear output power and utilizes nearly 50% of the entire radio transmit supply current. If the power efficiency of this stage is compromised, the entire radio power efficiency will suffer. To gain the best power efficiency for a single-ended 50-Ω type power amplifier, ground lead inductances must be extremely low <1 nH. It takes specialized packages and many bond wires to achieve these levels of inductance. For this reason, it makes the most technical and economic sense to put the power amplifier function in a dedicated IC. The RF and IF functions of the radio must next be partitioned. The IF SAW (surface acoustic wave) component is a bandpass filter at the IF frequency. It has nearly ideal frequency selectivity characteristics that are impractical to produce in an integrated circuit design. The SAW functions as the primary adjacent channel and interference blocking filter in the receive chain, and can be reused to remove spurious responses in the transmit chain. The RF and IF functions will meet at the SAW interface. For this reason, this interface is a natural place to consider dividing the functions, particularly considering the large size of the SAW. It would be difficult to maintain isolation around one RF and IF IC that had to interface to a SAW as well as other RF filters. Chip isolation is another consideration. The most sensitive RF function is the low noise amplifier in the receive chain; it must handle signals down to the microvolt level. Integrating all the RF and IF functions in a single IC would drive the pin count up and risk greater noise/interference since each additional pin has a potentially interfering signal. In addition, the practical board-level layout concerns and board component isolation (e.g., leakage into the SAW filter) make it more practical to divide the RF and IF functions into separate ICs.

C. Front-End IC Description

Fig. 5 shows the RF/IF IC. It consists of the receive side of a gain selectable LNA, downconvert mixer, and synthesizer. On the transmit side, there is the upconvert mixer, drive amplifier, and synthesizer. Receive and transmit circuitry can co-exist on the IC due to the half duplex operation of the radio where only one is on at a time. It was a goal of the design for cost reasons to try and eliminate the need for any external RF filters. For example, the image reject filter between the LNA and downconvert mixer on the receive side




Fig. 5. RF to IF converter.

could be eliminated if on-chip tuning networks could significantly attenuate image signals thereby preventing a noise figure degradation. On the transmit side, there is the same concern between the upconvert mixer and the drive amplifier, where image signals can cause linearity issues for the amplifier. There is the additional requirement on the transmit side of LO leakage being amplified without a filter between stages. Again, with on-chip tuning networks and isolation techniques to remove LO, both of these goals are met for this design. There is an additional filtering requirement at the IF port to utilize a single external SAW filter. This requires combining the Rx and Tx IF ports. Isolation and performance degradation are concerns, but through circuit and isolation techniques the goal of having a reusable IF SAW are achieved.

The cascaded gain of the Rx LNA/mixer is 25 db. Combined with a SSB (single sideband) NF of 3.7 db gives excellent sensitivity and effectively masks the impact of IF noise performance. The third-order intercept point of -13 dbm provides more than adequate linearity performance. The LNA is gain selectable with a low gain mode, which gives a cascaded gain of -5 db and input 1 db compression point (P1 db) of +2.5 dbm. This RF gain control is necessary as explained in the system requirements discussion. The transmit chain combination of mixer/drive amplifier has a cascaded gain of 25 db and an output P1 db of 4 dbm. The LO output level is typically <-50 dbm without external RF filtering. The transmit chain requires no external components for matching.

A 2.4-GHz programmable synthesizer is integrated on the RF part. Isolation and noise floor characteristics are major factors influencing this design. Differential high-frequency logic is used to combat isolation concerns and process 2.4-GHz signals. Once the logic signals are at low enough frequency, the integration and power efficiencies of CMOS logic are applied. In any phase locked loop system, within its loop bandwidth the synthesizer can dominate the close in phase noise performance. For this reason, this synthesizer

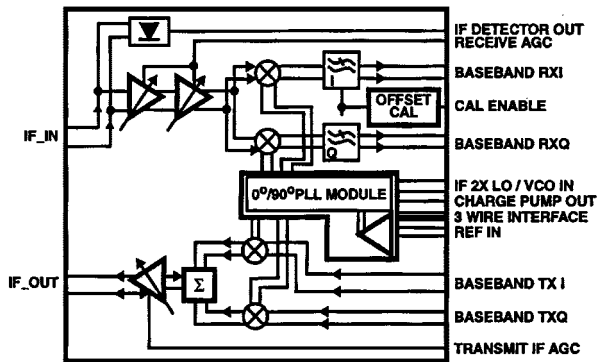


Fig. 6. IF/BB converter.

design incorporates circuitry to eliminate the so-called “dead zone.” A dead zone occurs when the finite charge pump response time of the synthesizer cannot keep up with very narrow correction pulses as directed by the closed loop system. When the charge pump is in a dead zone, the system is running open loop and the noise performance is degraded. Eliminating this zone of open loop operation has a positive effect on the overall phase jitter of the system and helps minimize the total charge pump current.

Referring to Fig. 6, the IF to baseband IC integrates all the Transmit and Receive functions necessary to process IF and analog baseband signals. The receiver has an IF energy detector, high-gain AGC amplifier, quadrature demodulator, low-pass filter, dc offset correction circuitry, and a 2X IF synthesizer. The transmit chain consists of a quadrature modulator and summer, IF AGC amplifier, and a 2X IF synthesizer. A 2X IF LO (e.g., Prism II with $IF = 374$ MHz used a 748-MHz LO input) is used so that digital divide by 2 circuitry, which gives excellent quadrature balance over a broad frequency range, can be implemented. The rest of the IF synthesizer is similar to the RF part. Again, a single IF port is shared by the Rx and Tx chains in order to reuse the external IF SAW filter. The entire Rx and Tx signal paths are differential to minimize isolation concerns.

The receive IF chain begins with the energy detector, which is used to determine if RF gain control is necessary. The output of this detector is sent to the baseband processor IC, where it is compared against a threshold and the appropriate digital control signal is provided to the RF part. The IF AGC amplifier follows next in the chain. This amplifier has a maximum high gain mode of 49 db of voltage gain at up to 600 MHz and a NF of 6.5 db. At the minimum gain setting, the AGC amplifier has -23 db of voltage gain and an input P1 db of -2 dbm (1.1 Vpp at 250- Ω impedance). The analog AGC switching is accomplished in <600 nanoseconds for a full scale change to within 1 db. The ac coupled quadrature demodulator is next. The requirement for <2 degrees of phase balance is achieved by the previously discussed divide by 2 circuitry up to a 1.2 GHz 2X LO input. A pair of two pole baseband low-pass filters used for anti-aliasing with 20 db of fixed voltage gain complete the signal path. In order to minimize the dc offset <10 mV after the filters, it is necessary to provide a dc offset correction loop. The dc offset

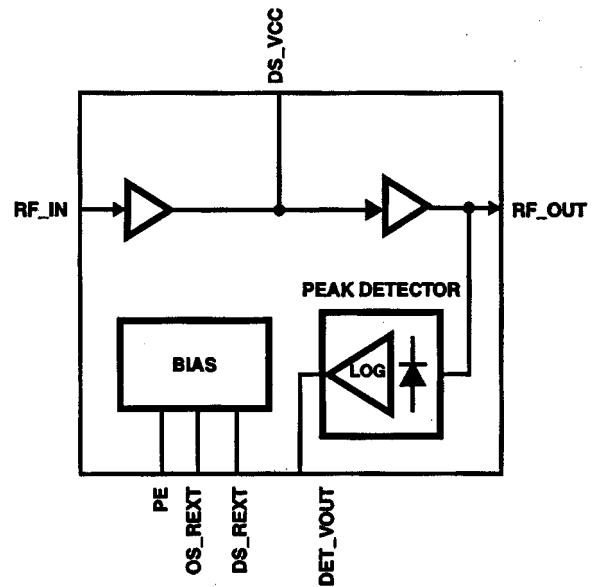


Fig. 7. Power amplifier.

correction can be run during transmit operation or when the radio is not actively receiving packets. The circuitry senses and corrects the dc offset by utilizing a digital loop, which has the advantage of infinite memory. The ac coupled AGC amplifier is shut down during this dc correction, while the demodulator and low-pass filters are active. The AGC loop is controlled by the baseband processor IC, which uses a sophisticated digital control loop.

The transmit IF section begins with the quadrature modulator, which performs in a manner similar to the receive demodulator. The modulator has the additional requirement of low carrier suppression <30 dbc, which means any dc offsets before or inherent in the modulator must be minimized. A transmit IF AGC is next. Power control is implemented in the Prism II radio by utilizing a power detector at the power amplifier output, which is fed back to an A/D on the baseband processor. After digitally processing this power signal, a D/A on the baseband processor will send the AGC level to the Tx IF AGC. Power control is important to the radio design to compensate for variation in the transmit chain from radio to radio, across channels, or over temperature. The Tx AGC output is typically a -10 dbm signal into a 250 Ω differential load. The AGC has over 70 db of attenuation range and maintains greater than 10 db compression margin above the output signal level in order to minimize any spectral regrowth.

The last of the three front-end ICs to discuss is the power amplifier shown in Fig. 7. This two-stage amplifier design typically delivers $+18$ dbm of linear power into a single ended 50- Ω load. The linearity of the amplifier as determined by its adjacent channel power rejection (ACPR) specification must be <-30 dbc to meet the 802.11 IEEE Standard. This amplifier requires no negative power supply like many GaAs amplifiers and has integrated bias and tuning networks on-chip. The integrated power detector function is accurate to within 1 db and allows precise power control of the radio and is key in maximizing power efficiency.

BASEBAND PROCESSOR




Fig. 8. Baseband processor.

D. Front-End IC Technology

Each of the front-end ICs presented particular challenges that influenced the choice of technology. It was decided for reasons particular to each IC to use what ended up being the same SiGe BiCMOS technology for all three ICs. The technology selected is a state-of-the-art process that includes triple-level metal, $0.35\ \mu\text{m}$ CMOS, trench isolation, and 60 GHz F_{max} bipolar transistors. Both the RF and PA parts require on-chip spiral inductors to minimize external matching components. These inductors require an excellent quality factor (Q), which is achieved by using last metal with trench isolation under the metal. Further, both the RF and PA parts must have excellent power efficiency for this portable application. The high F_t 's and F_{max} 's made available by the SiGe bipolar transistors allowed for RF power performance previously found only in GaAs technologies. The ultra-low-base resistance of SiGe transistors is another important consideration that gives a significant improvement in noise figure performance of the LNA and mixers. The RF and IF parts both required synthesizers which cannot be efficiently, from a power or area standpoint, built without CMOS logic. In addition, the large amount of functionality in the IF part including significant amounts of dc offset correction circuitry required additional CMOS circuitry. Therefore, the $0.35\text{-}\mu\text{m}$ CMOS allowed high levels of integration with low power. The overall front-end Tx and Rx signal chains had a nearly 50% reduction in power compared to the previous generation of product and technology. The high level of integration includes integrated synthesizers and eliminates external filters. This results in a nearly 50% reduction in the front-end IC costs to the end customers. Finally, these savings are possible while also increasing the overall front-end performance of this 11-Mb/s WLAN radio.

IV. BASEBAND PARTITIONING AND IMPLEMENTATION

This section presents the low-power consumption considerations for the digital BBP of the Prism II system, again capable of 1-, 2-, 5.5-, and 11-Mb/s direct sequence spread spectrum (DSSS) operations at 2.4 GHz. The DSSS spreading-code chip rate is 11 MHz. The BBP consists of the transmit section, carrier detect section, acquisition section and tracking section as shown in Fig. 8 [2]–[4]. Each of these sections contributes to the BBP power budget when operating in a WLAN. The 802.11 WLAN packet-based protocol establishes a BBP that cycles through well-defined states. When the radio is awake, the most common BBP state is listening for a valid signal. This is called carrier-detect processing, and figures heavily into the power

budget. If a valid signal is detected, the BBP transitions to the acquisition state. The acquisition state derives antenna diversity decision, timing offset, carrier frequency offset, a multipath-channel estimate, and optional equalizer design. The acquisition processing has the highest receive-mode power consumption per unit time, since its computation complexity is highest among all the receive states. This high power draw is mitigated by the short duration. It typically lasts less than $50\ \mu\text{s}$ and is a small overhead compared to the $800\text{-}\mu\text{s}$ duration of an 11-Mb/s 1-kB packet transmission. The tracking processing follows acquisition and recovers the transmitted information bits. The tracking processing performs channel-match-filtering, codeword correlation, optional equalization, and carrier-and-time tracking. Typically many stations reside in a cell, but only one is transmitting at a time. Since the BBP recovers data from every packet which it hears, it is important to minimize tracking state power draw. The transmit state is entered whenever a packet is available for transmission and the channel is deemed clear. The transmit state is high power draw because of the PA. The following subsections provide more detail about the BBP processing sections.

A. BBP Transmit Section

Several techniques can be used to make the BBP transmit section power efficient. The BBP transmit section presents a signal to the DACs, which is optionally generated either at I/Q baseband or at a digital IF. The I/Q baseband option consumes less power in the digital circuitry. On the other hand, the digital IF produces less distortion by eliminating dc offsets and I/Q amplitude-and-phase imbalances. The digital IF requires higher clock speeds, more-intensive digital transmit filtering and higher bit-precision DACs. Consequently, an I/Q baseband approach is more commonly used for power-sensitive WLAN applications. The digital IF will become more common as increasingly higher data rates are demanded by the marketplace. The higher data rates require the additional fidelity of a digital IF to minimize interference caused by the poorer I/Q balance of an analog based quadrature IF.

For DSSS modulation and I/Q baseband DACs, the transmit section tends to be the simplest component in the baseband processor, so the corresponding power drain is minimal. However, care must be taken to minimize the power-drain required for digital transmit filtering.

Transmit filter power drain can be reduced by using a ROM lookup of the transmit samples, rather than performing the power-intensive multiply-accumulate required by a convolution. This is accomplished by exploiting the fact that the DSSS codewords use BPSK and QPSK spreading chips. BPSK is a two phase modulation ($+1, -1$), and QPSK is a four phase modulation ($+1, +j, -1, -j$). Consequently, the in-phase (real) and quadrature (imaginary) baseband components can be described as impulse sequences containing only $+1$'s and -1 's. To achieve the desired spectral shaping, these impulse sequences are applied to a digital transmit filter. An FIR filter is commonly used, since an IIR prevents use of a ROM lookup. The FIR filter has a finite span. If the $+1$'s are

assigned 1's and the -1's are assigned 0's, the 1's and 0's in the FIR tap delay line act as an address. Consequently, the FIR output samples can be stored in a ROM and the chip impulse 1's and 0's used to address the appropriate output samples. Each polyphase up-sample FIR uses a different ROM table.

The PRISM II chipset uses an I/Q baseband approach and runs the 6-bit transmit DACs at $4 \times$ the chip rate. Four polyphase ROMs are used to synthesize the transmit signal at $4 \times$ the chip rate. The ROMs output 6-bit samples to the DACs. The same ROM is used for both the I and Q channels.

B. BBP Carrier Sense Section

When not in a sleep mode, the radio spends most of its time carrier-sense processing, so it is important to minimize BBP power draw here. Carrier-sense processing is used for both clear-channel-assessment (CCA) and acquisition activation. The protocol stipulates that a transmission is not initiated unless the channel is deemed clear. The channel is deemed clear if carrier sense is low. Carrier sense can be accomplished by jointly monitoring AGC behavior and Barker codeword correlation. The preamble spreading sequence is an 11-chip Barker code (+1, -1, +1, +1, -1, +1, +1, +1, -1, -1, -1). This same spreading sequence is used for the data payload portion of 1- and 2-Mb/s packets.

In an AGC'd receiver, the AGC can be used as part of the carrier sense and is a low-power process. A significant, sudden increase in received signal power denotes the probable onset of a packet. The application may or may not care about very weak signals, say less than a 3-dB SNR. If these weak signals can be ignored, a great deal of power savings is realized by shutting down digital processing until the AGC detects a signal above the 3-dB threshold. If very weak signal reception is a goal, Barker correlation can be used for carrier detection. The 1-Mb/s signal is capable of operating at an SNR of 0 dB, so here Barker correlation on the noise floor is important.

Barker correlation is used to detect weak signals and to qualify strong signals as valid or not valid. Barker correlation is accomplished with low-power draw because it can be accomplished with merely adds and subtracts. Microwave-oven radiation is an example of an invalid strong signal. The preamble consists of biphase-modulated 11-chip Barker codewords. Barker codes have a nearly impulsive autocorrelation function. A signal containing a repetition of Barker codes is deemed valid by the receiver. This repetitious characteristic is distinguishable even in the presence of heavy multipath. The Barker code provides more than 10 dB of spreading gain.

C. BBP Acquisition Section

If a valid signal is detected, acquisition processing derives the antenna diversity decision, the timing offset, the carrier frequency offset, multipath channel estimation, and optional equalizer design. In this state, power drain can be minimized during antenna selection and channel-impulse-response estimation.

The antenna diversity decision can be made by either selecting the antenna with the most power or by choosing the antenna with the least amount of multipath. The AGC is used




Fig. 9. RAKE receiver for 1- and 2-Mb/s signal reception.




Fig. 10. RAKE receiver for 5.5- and 11-Mb/s signal reception.

to measure the amount of power on an antenna. Choosing the antenna with the least amount of multipath is more computationally intensive and timeline consuming, since a multipath estimate is required. Choosing the antenna with the most power is simple and results in the lowest power drain.

The timing-offset is eliminated by centering codeword detection on the multipath profile. The multipath profile is estimated via Barker correlation. The frequency offset is eliminated by measuring the amount of signal rotation that occurs on successive Barker codewords.

The multipath estimate is used to design the channel matched filter (CMF). Ideally the CMF is the time-reversed, conjugate of the channel impulse response (CIR). There are low-complexity multipath-estimation algorithms and high-complexity algorithms. The high-complexity algorithms provide higher fidelity. The simplest way to estimate the channel is to coherently average the Barker correlator's output over multiple received codewords. The frequency/phase estimate establishes the coherence. The Barker codeword does not have a truly impulsive autocorrelation function. Nevertheless, it oft-times provides sufficient fidelity for most applications. The biphase modulation of the Barker word during the preamble helps minimize the correlation sidelobe distortion through sidelobe modulation.

D. BBP Tracking Section

The power drain during tracking can be minimized in the channel-matched-filtering stage and codeword-correlation stage.

1) *RAKE Receiver*: The heart of a high-performance 802.11 DSSS modem is the RAKE receiver. The RAKE receiver approaches an optimal matched-filter receiver for communications in multipath and noise [5]–[7]. While multiple canonical forms exist, the most efficient architecture for 802.11 high-rate transmissions requires that the codeword correlator follow the channel matched filter (CMF) as shown in Figs. 9 and 10. In most references, the codeword correlator is shown in front of the channel matched filter (CMF). The opposite is preferred here because of the large number of codewords. These are the two computationally dominate blocks in the RAKE. The CMF consists of a complex tap for each multipath finger. Each multipath finger has a unique

amplitude and phase. Usually, only a few bits of precision are required to model each channel-finger in the CMF.

To save power, only the strongest multipath fingers should be tracked. This reduces the number of complex multiplies required when convolving the signal with the CMF. PRISM II uses a channel matched filter with a resolution of two samples/chip. However, the output of the CMF is decimated to the chip rate. Power savings is realized by performing multiply accumulates only at the chip rate.

For 1- and 2-Mb/s reception, codeword correlation is accomplished using a single codeword—the Barker code. For 11-Mb/s reception, correlation is required for 64 different codewords. At 11 Mb/s, 6 bits are mapped to 1-of-64 codewords and 2 bits are mapped to 1-of-4 quadriphase rotations of the selected codeword. It is important to make the correlation for 64 codewords efficient. The next section explains how this can be accomplished at 11 Mb/s.

2) *Fast CCK Transform*: The DSSS spreading codes for 5.5 and 11 Mb/s use eight chips per codeword, rather than the 11 chips of the 1- and 2-Mb/s modes. All data rates use a chip rate of 11 MHz. The eight chips are taken from a class of codes called complementary code keying (CCK). The CCK codes are defined in [8] and are a type of Walsh code generalized to complex elements. Just like Walsh codes, CCK can be decoded using a modified fast Walsh transform (FWT).

Ideal codes for use in multipath would have impulsive autocorrelation functions and zero cross-correlation functions. Since this cannot be realized in practice, the idea behind complementary codes is to make the composite sum of individual autocorrelation functions combine to an impulse. This property is called complementary. In [8], the codes are formed by first making a kernel which is one complementary pair from which all other complementary sequences can be derived. The kernel is formed using Golay's rule for length expansion. As shown in [8], if we start with the length 4 sequence A_1B_1 , where $A_1 = [1 \ 1]$ and $B_1 = [1 \ -1]$, a recursive rule for expansion is: $A_n = A_{n-1}B_{n-1}$ and $B_n = A_{n-1}B'_{n-1}$ ($B' = -B$). Thus, the length 8 sequence is $[1 \ 1 \ 1 \ -1 \ 1 \ 1 \ -1 \ 1]$.

Using these codes, the next step is to find $\log_2 N$ orthogonal subsets (where N is the length of the code) for this code. These subsets are formed from the even ordered single elements, pairs and quads. According to [9], each subset can be given a different phase without disturbing the complementary code characteristics. The following equation represents this set of codes based on the kernel given above for $N = 8$. Each phase variable is four-valued. This equation derives all the codewords used in the high-rate extension of 802.11 at 2.4 GHz. Sixty-four base codewords exist, and the others are just quadriphase rotations of the base 64 (see (4.1) at the bottom of the page).

$$c = \left\{ e^{j(\varphi_1+\varphi_2+\varphi_3+\varphi_4)}, e^{j(\varphi_1+\varphi_3+\varphi_4)}, e^{j(\varphi_1+\varphi_2+\varphi_4)}, -e^{j(\varphi_1+\varphi_4)}, e^{j(\varphi_1+\varphi_2+\varphi_3)}, e^{j(\varphi_1+\varphi_3)}, -e^{j(\varphi_1+\varphi_2)}, e^{j(\varphi_1)} \right\} \quad (4.1)$$

Basic Fast Walsh Block (BFWB)




Fig. 11. Modified fast Walsh transform: basic fast Walsh (BFWB).

Based on this analysis, the minimum distance between two different complementary codes is $N/2$ symbols. Therefore, it is possible to correct $N/4 - 1$ symbol errors. If M phases are possible, this yields a minimum Euclidean distance of

$$d_{\min} = \sqrt{\frac{N}{2} \left\| 1 - \exp\left(j\frac{2\pi}{M}\right) \right\|}. \quad (4.2)$$

The four phase variables each take on values of $[0, \pi/2, \pi, 3\pi/2]$, and there are 256 possible eight-chip codes. These codes have an inherent “Walsh” type structure that allow a simple butterfly implementation of the decoder. Although it is possible to squeeze a few more complementary codes out of this eight-chip set, the rest of the codes cannot be decoded with the modified fast Walsh transform.

Fig. 11 shows the basic fast Walsh block, which brings in eight chips of soft decision data shown here by x_0, x_1, \dots, x_7 , and produces 16 possible correlations for given values of φ_1 and φ_2 . Each individual chip takes on the value 1, -1 , j or $-j$. Hence, the fast transform can be implemented using only adds, subtracts, and real-imaginary data swapping. For additional information on signaling and coding waveforms, see [10]–[25].

Fig. 12 shows all 256 possible correlator outputs.

E. Additional BBP Power Considerations

The WLAN protocol is amenable to various power saving modes. The system can be configured to allow the radio to be dormant for prescribed periods of time, with the radio waking up at intervals. The time to resume stable operation after being in “sleep” mode is an important consideration. While simply turning off the supply voltage to the baseband would

8 chip Complementary Code Decoder




Fig. 12. Modified fast Walsh transform.

seem to be the ultimate power savings mode, it is in fact impractical. The MAC (media access controller) chip must remain active to restart the radio, as well as control TX and RX operations. Interface lines between a powered-up chip and an unpowered device are unwieldy, and the loss of programmed register settings and learned behavior of the baseband makes powering down the device undesirable.

The mixed signal blocks were designed with two levels of standby. The first level maintains bias networks which allows for return to stable operation within $2 \mu\text{s}$ of appropriate signal assertion. The second level allows complete shutdown to reduce standby currents to CMOS leakage levels. Several milliseconds are required to return to stable operation.

The Prism II architecture has two separate blocks, one for 1 and 2 Mb/s and a second for 5.5 and 11 Mb/s. The progressive scaling in clock speeds required by each block naturally suggests multiple clock domains for the various receive and transmit functions that maximize power. However, this conflicts with the desire for fast and uncomplicated gate synthesis and static timing analysis as provided by a fully synchronous, single-clock-domain design. The compromise Prism II solution is to use five clock domains: one for the serial interface, two for transmit, and two for receive. On the transmit side, all the digital logic runs at the 11-MHz DSSS chip rate; the output FIR and DACs run in a 44-MHz domain. Neither of these two clocks runs unless transmit operation is active. On the receive side, the block that calculates the channel impulse response (CIR) runs on a 44-MHz domain to minimize delay. This section only runs when an initial carrier-detect condition is met. It is turned off after CIR estimation is complete. Even within this 44-MHz block, data paths are zeroed until specifically called for. The remainder of the receiver runs on a single 22-MHz clock. The interpolation buffer, downconverter and channel-matched-filter have data paths zeroed until needed, minimizing power draw when inactive. The various data paths and state machines for 1-, 2-, 5.5-, and 11-Mb/s sections do not run unless needed.

V. MAC PARTITIONING AND IMPLEMENTATION

In order to provide access control, addressing, sequencing, checking, and delimiting functions for delivery of PDUs between peer LLC entities over the wireless physical layer, the IEEE 802.11 MAC provides several additional functions, including acknowledgment, association, authentication, distribution, power management, and privacy.

A. Service Sets

The set of stations “connected” to a WLAN is often different from the (dynamically changing) set of stations that are able to exchange frames over the wireless medium. This requires the IEEE 802.11 MAC to provide a means of identifying those stations logically connected to the network, and to tolerate *hidden stations*, which receive an indeterminate subset of the relevant transmissions, and whose transmissions are received by a different, indeterminate subset of relevant stations. A *basic service set (BSS)* is the set of IEEE 802.11 stations that comprise a logical network. The process by which a station joins a BSS is known as *association*, and involves the exchange of management frames to ensure that the station implements an acceptable set of MAC and PHY functions and to establish a common set of operational parameters for use by all stations in the BSS. IEEE 802.11 MAC headers include a *BSS identifier (BSSID)*, so stations can detect and discard (valid) frames received from stations in other BSSs operating in the same volume of space.

To provide wireless coverage for an area comparable to the physical extent of a wired LAN, IEEE 802.11 uses an *extended service set (ESS)*. An ESS is a set of BSSs that appear to be a single BSS to the LLC sublayer at any station associated with any of those BSSs. An *infrastructure BSS* is capable of being a member of an ESS, and includes an *access point (AP)*, which is a station that provides association, authentication, and distribution services for the BSS. The IEEE 802.11 MAC includes a *scanning* function, which permits stations to locate nearby BSSs. Scanning may be either passive, where the station listens for the reception of *Beacon* frames which are transmitted periodically by each AP; or active, where the station broadcasts *Probe Request* frames, which elicit a *Probe Response* frame from each AP receiving the probe request.

The AP may also provide a *portal* function, typically implemented as an IEEE 802.1d MAC bridge, to permit communication between stations on the wired infrastructure and stations associated with the BSS. Distribution services are used to forward frames between associated stations over the wireless medium as well as to transfer frames between APs in the ESS. Stations may transfer their association among BSSs of an ESS while retaining network connectivity by performing a *reassociation*. IEEE 802.11 also allows an *independent BSS (IBSS)*, often referred to as an “*ad hoc* network.” An IBSS may be convenient for communicating among a dynamic or temporary collection of stations, but lacks association and distribution services, and can only provide a subset of the power management and frame delivery mechanisms available in an infrastructure BSS. Because the functionality of an IBSS is a subset of a BSS, it will not be discussed further.




Fig. 13. IEEE 802.11 frame format.

B. MAC Frames

The use of a physically open network medium, operating in a radio frequency band that is shared with other (non-WLAN) usage, complicates the accurate delimitation of frame boundaries. The general format of IEEE 802.11 frames is shown in Fig. 13. The beginning of each frame is identified by a starting delimiter with a bit pattern which is uniquely distinguishable from the preceding PHY preamble. Following the starting delimiter (SFD) is a *Physical Layer Convergence Protocol (PLCP)* header that conveys the frame length as well as PHY-specific signaling information. The PLCP header error check (HEC) uses a 16-bit CRC. The end of a frame is determined from the time following the SFD needed to transfer the number of octets indicated in the PLCP length field at the data rate indicated in the PLCP signal field. Use of a CRC-protected length instead of an ending delimiter allows an accurate determination of the proper ending time, even if the incoming signal is lost during frame reception. This is needed for proper operation of the CSMA/CA arbitration protocol. Accurate length indication also permits the 32-bit CRC used for the MAC FCS to be checked at the proper point in the received data stream, thereby achieving an erroneous frame acceptance rate of less than 2.3×10^{-10} .

Following the PLCP header is the MAC header, which contains frame control, type, address, and sequence control information, as well as a duration field that indicates the amount of time to complete the next stage in the current *frame exchange sequence*. There are three, and in some cases four, address fields in IEEE 802.11 data and management frames because the dynamic topology of a WLAN requires an explicit distinction between the SA (source address, identifying the MAC entity that initiated transfer of the MSDU), DA (destination address, identifying the MAC entity or entities intended to be the final recipient(s) of the MSDU), TA (transmitter address, identifying the station that transmitted the MPDU onto the wireless medium), and the RA (receiver address, identifying the station(s) intended to receive the MPDU from the wireless medium). On a wired LAN, only two address fields are needed because $TA = SA$ and $RA = DA$.

C. Frame Exchange Sequences

The BER achieved by wired LAN physical layers is generally better than 10^{-7} , and does not vary significantly over (short periods of) time, so wired LAN protocols, and higher layer protocols intended for use over LANs, are intended to operate in the presence of frame loss rates around 0.1%. The

BER achieved by wireless physical layers varies dynamically over both time and space, and often is two orders of magnitude poorer than wired physical layers. LANs are not expected to provide reliable delivery, but most LLC, network, and transport layer protocols used for reliable and/or acknowledged delivery are “success oriented,” favoring throughput over timely retransmission. System performance and perceived usability suffer severely if higher-layer recovery mechanisms intended for use on 0.1% of transmitted frames need to be invoked for 10% of those frames. The IEEE 802.11 MAC acknowledges successful reception of unicast frames, and automatically retransmits unacknowledged frames (with separate retry counts for long and short frames, each controllable by the management entity), thereby significantly reducing the frame loss rate visible to LLC and higher layer entities. Transmission of the acknowledgment control frame follow the end of the frame being acknowledged by a short inter-frame space (SIFS), which is two slot times shorter than the DCF inter-frame space (DIFS) used by stations contending for the right to transmit. This causes the entire frame exchange sequence to appear as a single instance of network activity to deferring stations.

The use of acknowledgment and retransmission within the MAC sublayer introduces the possibility of *frame duplication*. Frame duplication occurs when a unicast data or management frame is received successfully at the RA station, but the acknowledgment control frame is not received successfully by the TA station, thereby causing an unnecessary retransmission. Since higher protocol layers are not required to tolerate duplicated frames, the IEEE 802.11 MAC header includes a retry bit and a sequence control field that combine to permit receiving stations to detect and discard duplicates.

When a WLAN is operating in an environment with large amounts of RF interference, or a large population of stations contending for the medium, retransmissions of unacknowledged frames may result in inefficient use of time on the wireless medium, which tends to be the scarcest resource within the BSS. Therefore, the IEEE 802.11 MAC allows a frame exchange sequence for MSDUs longer than a specified threshold to begin with the exchange of a small request to send (RTS) and clear to send (CTS) control frames between the TA and RA stations. A successful RTS-CTS exchange establishes that the pair of stations are not hidden from each other, and activates the virtual carrier sense mechanism at other stations surrounding both the TA and RA stations, thereby reducing substantially the risk of wasting time transmitting the subsequent long frame.

D. Coordination Function

The IEEE 802.11 MAC uses a *distributed coordination function (DCF)* to arbitrate among stations for the right to transmit onto the wireless medium. The DCF uses carrier sense multiple access with collision avoidance (CSMA/CA) with binary exponential backoff (BEB). Collision avoidance is used instead of the more common collision detection (CSMA/CD) because useful collision detection requires full connectivity. With a wireless physical layer, the collision may prevent frame reception at the intended recipient station while




Fig. 14. RTS/CTS/data/ACK frame exchange and NAV setting.

the two (or more) transmitting stations are unable to receive each other's signals. To provide further protection from destructive interference due to uncoordinated transmissions by hidden stations, the carrier sensing employs both the detection of physical activity on the wireless medium, using the PHY layer's CCA function, and a virtual carrier sense mechanism known as the *network allocation vector (NAV)*.

The NAV can be considered to be a decrementing timer that is nonzero during periods when a frame exchange sequence is (believed to be) in progress. IEEE 802.11 MAC headers include a duration field, which contains the time from the end of this frame until the end of the current frame exchange sequence. Any station receiving an error-free frame that is not addressed to itself updates its NAV to contain the greater of the present NAV value or the value of the newly received duration field. Whenever the NAV is nonzero, the DCF considers the medium to be busy, independent of the state of CCA from the PHY. A simple example of NAV operation is shown in Fig. 14. Note that the NAV has the same ending time whether set from the RTS frame, the CTS frame, or both. This reduces the probability of destructive interference at the addressed recipient station, because other stations in the vicinity of the recipient are likely to receive the CTS frame transmitted by that station, even if they are not able to detect the RTS frame and subsequent data frame transmitted by the SA station. Use of the NAV also reduces interference between stations in overlapping BSSs because NAV updates from received frames are independent of BSSID.

The DCF includes three mechanisms that combine to prevent the well-known "capture effect" problem that occurs on most CSMA protocols with BEB, including IEEE 802.3, under heavy load conditions. First, stations only decrement their back off timers when the medium (both CCA and NAV) is not busy, so the relative queue positions of (most) contending stations is unchanged during periods of deferral to a busy medium. Second, all stations execute a backoff procedure after successful transmission, which randomizes their reentry into the set of stations contending for the medium. Third, separate retry counts are used to determine when to reset the contention window size and when to abandon transmission attempts for a given frame.

The IEEE 802.11 MAC provides an optional *point coordination function (PCF)* that can be used in infrastructure BSSs. When active, the point coordinator periodically obtains control of the medium using the DCF, then controls transmissions during a *contention free period*, of predefined maximum length, by initiating transmissions and polling associated

stations. During the contention free period, utilization of the wireless medium is much more efficient, primarily because no time is lost due to backoffs nor collisions, but also because all frames can be separated by SIFS intervals, and acknowledgment can be piggybacked on subsequent data or poll frames, saving the time to transmit a separate ACK frame. The PCF is well suited to providing time-bounded, isochronous, reserved bandwidth, and/or other policy-based quality of service (QoS) functions within a BSS. Since the PCF operates on top of the DCF, and contention free periods alternate with contention periods, IEEE 802.11 does not define the specific services or policies provided by a point coordinator.

E. Power Management

Power consumption is not an issue in the design of wired LANs, because utility power is assumed to be available at every wired LAN connection point. Stations on WLANs commonly need to operate from battery power, which makes the reduction of power consumption an important design consideration. Most of the opportunity to reduce power consumption by circuit implementation techniques exists at the PHY layer, because the MAC implementation is composed of a set of communicating finite state machines, typically implemented as software or firmware. Most of these state machines perform no state transitions when the station is not engaged in a transmission or reception attempt. However, the system level power consumption can be reduced significantly by including appropriate power management functions within the MAC protocol.

A side effect of this nonconcern about power consumption in the design of wired LANs is that IEEE 802.2 LLC, and most higher layer protocols, implicitly assume that the receivers at all stations are enabled continuously. The IEEE 802.11 MAC protocol permits major power savings by allowing the PHY receiver, not just the PHY transmitter, to be turned off for significant periods without the station appearing to be disconnected from the LAN. Except during short bursts, such as while writing to remote file servers or printers, the transmitter at a WLAN station is active less than 2% of the time. Therefore, over the life of a typical battery charge, more than 90% of the power used by a station is consumed in the PHY receiver circuitry. During periods when a station is active but not performing LAN transfers, the entire PHY can be turned off as much as 95% of the time, and for low-intensity LAN communication the receiver can be off 60% to 75% of the time. This can extend battery life far beyond what is achievable strictly through power conserving circuit design, because the power consumption of an enabled but quiescent receiver cannot be reduced to zero. It has been shown that stations using power save mode, associated in a point coordinated BSS, can reduce power consumption, averaged over 1-min intervals, by as much as 97.5% relative to stations using the same PHY implementation but active continuously.

In an IEEE 802.11 BSS, the AP buffers broadcast and multicast frames, as well as unicast frames addressed to stations that have indicated to the AP that they are operating in power save (PS) mode. Each Beacon frame transmitted by the AP includes a *Traffic Information Map (TIM or DTIM)*, a compressed bitmap that indicates which associated stations have

buffered traffic. The *DTIM interval* BSS operational parameter defines the number of beacons containing TIMs between beacons containing DTIMs (typically zero). Stations in PS mode activate their receivers once each listen interval, at the target beacon transmission time (TBTT) for beacons containing DTIMs. Stations in PS mode may also activate their receivers at TBTTs for beacons containing TIMs (if any). Buffered broadcast and multicast frames are transmitted by the AP immediately following each beacon that contains a DTIM, thereby allowing these group-addressed frames to be received by all associated stations. Following the reception of any TIM, or after the group-addressed frames that follow a DTIM, PS stations that detect the absence of buffered traffic can disable their receivers until TBTT for the next DTIM, with no risk of missing a relevant frame.

If a PS station receives a TIM or DTIM that indicates the existence of buffered frames, that station may elect to remain active, and request delivery of a buffered frame by sending a PS-Poll control frame to the AP. In response to each PS-Poll, the AP transmits the oldest buffered unicast frame for the TA of the PS-Poll. The “More Data” bit in the frame control field of this frame is set if additional unicast frames are buffered for this station, which allows the PS station to determine whether there is benefit in remaining active and sending additional PS-Poll frames. The AP ages buffered frames, and discards them if they remain undelivered for an excessive period.

There is an increase in transaction latency when communicating with stations using PS mode. This added latency is bounded, and can be controlled by changing the values used for the beacon interval and the DTIM interval to avoid causing timeouts at higher protocol layers. The AP indicates the existence of buffered, unicast frames, but only transmits those frames in response to requests from the PS station. This permits local decision making at the remote (battery operated) station about how long to leave its receiver enabled, when to request delivery of buffered frames, and whether to switch to active mode temporarily during bursts of heavier network activity involving that station.

F. Security and Privacy

Wired LANs assume the owner of the LAN has physical control over the network, and can take appropriate measures to secure the physical medium from unauthorized access. Under this assumption, receipt of a well-formed frame from the wired medium is *prima facie* evidence that the sending station is authorized to communicate on the network. The wireless medium is physically open, with no easily enforced barriers to unauthorized transmission of frames that appear to belong on a particular network, nor to the unauthorized reception of frames (“eavesdropping”). The IEEE 802.11 MAC includes an *authentication* function to establish the identity of a station prior to allowing the that station to associate with a BSS or to exchange data frames with other stations. There is also an optional *wired equivalent privacy (WEP)* function that performs encryption of MAC payload (MSDU) data for transfer over the wireless medium. The authentication and WEP functions are do not provide security, which is a function provided at higher layers of the protocol stack. These functions are intended to allow the deployment and use of

WLANs in equivalent environments as wired LANs by providing a qualitatively similar levels of connection control and communication privacy as inherently exist on wired LANs.

VI. SUMMARY AND FUTURE EXPECTATIONS

As is evident in this paper, WLANs have a significant number of technical issues that must be considered when developing both the system and circuit level solutions. Among these is low-power operation. The biggest impact on minimizing power is in the total system architecture where clever use of sleep protocols and intelligent listening for valid signals is utilized. At the circuit level, specific performance requirements must be met, and power is frequently a design tradeoff. The general approach is to assure that parameters such as range (receiver noise figure, PA output power) and multipath (linearity) are met with the most power-efficient approaches possible. This may mean that power is not the primary consideration in all cases. As a result, the protocol is designed to provide power management so that critical performance requirements can be met without fear of significant power drain.

Future WLAN networks will have similar environmental constraints, but will also need to run at higher data rates. At the same time, technology will advance to combat and offset the potential power increases. New signaling schemes that reduce the peak to average power requirements and permit reduced PA backoff are in development. Such signaling schemes may also be more immune to nonlinear channel effects, which will permit higher efficiency PA approaches and, hence, have significant power impact. More efficient radio architectures such as zero IF or image rejection/low IF approaches will minimize the amount of off-chip I/O and drive requirements. Finally, IC process technology will continue to advance via Moores law for the next several years. Smaller physical dimensions will permit faster device speed at much lower current. This will have an impact from the front-end LNA all the way through to the host interface.

With these advances and others not yet envisioned, highly portable, high-speed and power-efficient wireless data networks will continue to evolve.

REFERENCES

- [1] *IEEE Local and Metropolitan Area Networks: Overview and Architecture*, IEEE Standard 802-1990, 1990.
- [2] M. Webster *et al.*, Harris/Lucent T/Gb compromise CCK (11 Mb/s) proposal, July 1998. doc.: IEEE P802.11-98/246.
- [3] M. Webster, Multipath issues and architectures, Jan. 1998. doc.: IEEE P802.11-98/37.
- [4] ———, DFE Packet error rate minimization using precursor sliding, Mar. 1998. doc.: IEEE P802.11-98/115.
- [5] R. van Nee, Delay spread requirements for wireless networks in the 2.4 GHz and 5 GHz bands, Nov. 1997. doc.: IEEE P802.11-97/125.
- [6] H. Hashemi, “The indoor propagation channel,” *Proc. IEEE*, vol. 81, pp. 943–968, July 1993.
- [7] J. G. Proakis, *Digital Communications*, 3rd ed. New York: McGraw-Hill, 1995.
- [8] R. van Nee, “OFDM codes for peak-to-average power reduction and error correction,” in *IEEE Global Telecommunications Conf.*, Nov. 18–22, 1996, pp. 740–744.
- [9] R. L. Frank, “Polyphase complementary codes,” in *IEEE Trans. Inform. Theory*, Nov. 1980, pp. 641–647.
- [10] M. J. E. Golay, “Complementary series,” in *IRE Trans. Inform. Theory*, Apr. 1961, pp. 82–87.

- [11] C.-C. Tseng and C. L. Liu, "Complementary sets of sequences," in *IEEE Trans. Inform. Theory*, Sept. 1972, pp. 644–652.
- [12] R. Sivaswamy, "Multiphase complementary codes," in *IEEE Trans. Inform. Theory*, Sept. 1978, pp. 546–552.
- [13] F. Kretschmer and K. Gerlach, "Radar Waveforms Derived from Orthogonal Matrices," NRL Rep. 9080, Feb. 1989.
- [14] K. Gerlach and G. Kretschmer, "General forms and properties of zero crosscorrelation radar waveforms," in *IEEE Trans. Aerosp. Electron. Syst.*, Jan. 1992, pp. 98–103.
- [15] R. Sivaswamy, "Digital and analog subcomplementary sequences for pulse compression," in *IEEE Trans. Aerosp. Electron. Systems*, Mar. 1978, pp. 343–350.
- [16] F. F. Kretschmer Jr. and B. L. Lewis, "Doppler properties of polyphase coded pulse compression waveforms," in *IEEE Trans. Aerosp. Electron. Syst.*, July 1983, pp. 521–531.
- [17] G. Weathers and E. M. Holliday, "Group-complementary array coding for radar clutter and rejection," *IEEE Trans. Aerosp. Electron. Syst.*, pp. 369–379, May 1983.
- [18] A. Ojha and D. Koch, "Performance analysis of complementary coded radar signals in an AWGN environment," in *IEEE Proc. Southeastcon*, Apr. 8–10, 1991, pp. 842–846.
- [19] K. V. Prasad and M. Darnell, "Data transmission using complementary sequence sets," in *IEEE 5th Int. Conf. HF Radio Systems and Techniques*, July 22–25, pp. 222–226.
- [20] B. Popovic, "Synthesis of power efficient multitone signals with flat amplitude spectrum," in *IEEE Trans. Commun.*, July 1991, pp. 1031–1033.
- [21] C. Tellambura, Y. J. Guo, and S. K. Barton, "Channel estimation using aperiodic binary sequences," in *IEEE Commun. Lett.*, May 1998, pp. 140–142.
- [22] *IEEE*, IEEE Standard 802.11-1997.
- [23] C. Heegard *et al.*, Performance of PBCC and CCK, Sept. 1998. IEEE 802.11-98/304.
- [24] T. S. Rappaport, *Wireless Communications—Principles and Practice*. Englewood Cliffs, NJ: Prentice-Hall, 1996.
- [25] K. Halford *et al.*, Complementary code keying for rake-based indoor wireless communication, in IEEE ISCAS'99, Orlando, FL, to be published.



Brent A. Myers (Member, IEEE) received the B.S. degree from Purdue University, West Lafayette, IN, in 1979, the M.S. degree from Virginia Polytechnic Institute, Blacksburg, in 1982, and the Ph.D. degree from the Florida Institute of Technology, Melbourne, in 1994, all in electrical engineering.

From 1979 to 1985, he worked at the Mobile Communications Division of General Electric, Lynchburg, VA, where he was involved in analog IC design for commercial and emergency mobile communication systems, and the AMPS cellular system. In 1985, he joined Harris Semiconductor, Melbourne, FL, where he was a Design Engineer for a wide variety of application-specific circuits, including IR imaging networks, tape driver equalizers, and analog baseband processors for telecommunication and wireless systems. He is presently Senior Scientist and Technical Leader of the mixed-signal product development group at Intersil Corporation, Melbourne, FL, where he is involved in circuit design for wireless LAN and high-speed wireline modems. He is an Adjunct Professor at the Florida Institute of Technology in the area of analog and mixed-signal design. He has eight patents in the areas of signal processing and analog circuit techniques.



J. Bartow Willingham (Member, IEEE) received the B.E.E.E. and M.S.E.E. degrees from the Georgia Institute of Technology, Atlanta, in 1976 and 1977, respectively.

He joined Harris Government Communication Systems Division in 1977, where he developed modems and communication terminals for the Department of Defense and NASA. He has been with Intersil Corporation, Melbourne, FL, since 1998, where he is currently a Wireless System Engineer.



Patrick Landy was born in New York, NY, on September 12, 1963. He received the B.S.E.E. degree from the University of South Florida, Tampa, in 1985 and the M.S.E.E. degree from the Florida Institute of Technology, Melbourne, in 1989.

Since 1986, he has been an IC Circuit Designer with Intersil Corporation, Melbourne, FL. He is currently a Senior Principal Engineer in the mixed-signal product development group, where he is involved in analog and RF IC design for wireless LAN networks.




Mark A. Webster (Member, IEEE) received the B.S. and M.S. degrees in electrical engineering from the University of Michigan, Ann Arbor, in 1980 and 1981, respectively.

From 1981 to 1997, he was with the Government Communication Systems Division, Harris Corporation, Palm Bay, FL. Since 1997, he has been with Intersil Corporation, Melbourne, FL. His interests include adaptive equalization and signal processing algorithms for both wireline and wireless communications.



Perry Frogge was born in Melbourne, FL, in 1958. He received the B.S.E.E. degree from the University of Central Florida, Orlando, in 1980.

Since 1980, he has been a Digital CMOS IC Designer at Intersil Corporation, Melbourne. He started with 80C86, 80C286 processor and peripheral design, followed by work in DSP, and for the past five years has done WLAN, baseband processor design. He holds several patents and application note awards.



Michael Fischer (Member, IEEE) received the B.S. degree in geology from Beloit College, Beloit, WI, in 1975.

From 1972 to 1974, he was a Design Coordinator for the MAPS Project at the University of Kansas, creating one of the first raster graphics systems. From 1975 through 1982, he was Vice-President at CHILD Incorporated, leading the design of display generators, image processors, and programmable controllers. From 1983 to 1984, he was system Architect at Syte Information Technology, San Diego, CA, defining graphics workstations. From 1985 to 1989, he was Principal Architect at Datapoint, San Antonio, TX, developing symmetric multiprocessor servers and a LAN with dynamically variable data rates. In 1990, he became President of CHILD Systems, where he developed custom wired and wireless networks for medical and industrial applications. Since 1994, he has focused on wireless LANs, first at Digital Ocean, then at CHOICE Microsystems after its 1998 spinoff. He is currently CTO of CHOICE-Intersil, a subsidiary of Intersil Corporation, Melbourne, FL, where he develops WLAN MAC protocols and controllers.

Mr. Fischer has been a voting member of, and major contributor to, the IEEE 802.11 WLAN standards group since 1993, and was State Machine Editor for IEEE Standard 802.11-1997 and ISO/IEC 8802.11:1999. He is a Member of the IEEE Computer and Communications Societies and the ACM.