

Design and Implementation of an All-CMOS 802.11a Wireless LAN Chipset

Teresa H. Meng, Stanford University

Bill McFarland, David Su, and John Thomson, Atheros Communications

ABSTRACT

The tremendous growth in wireless LANs has generated interest in technologies that provide higher data rates and greater system capacities. The IEEE 802.11a standard [1], based on coded OFDM modulation [2], provides nearly five times the data rate and at least 20 times the overall system capacity compared to the incumbent 802.11b wireless LAN systems [3]. This article describes the design challenges and circuit implementation of a two-chip set that forms a complete 802.11a solution in 0.25 μm CMOS technology. Wherever possible, sophisticated digital signal processing techniques are used to compensate for possible analog impairments associated with integrating RF circuitry in a CMOS technology. The analog portion of the chip set implements a 5 GHz transceiver comprising all the necessary RF and analog circuits of the 802.11a standard integrated on a single chip. Some features of this IC include 22 dBm peak transmitted power, 8 dB overall receive-chain noise figure, and -112 dBc/Hz synthesizer phase noise at 1 MHz frequency offset. The digital portion of the chip set, the baseband and MAC processor, contains dual ADCs/DACs and all the digital circuits for synchronization, detection, and 802.11 MAC layer data processing. This IC delivers up to 54 Mb/s in a 20 MHz channel according to the 802.11a standard, and includes proprietary modes supporting up to 108 Mb/s in a 40 MHz channel.

INTRODUCTION

Since 1999, the wireless LAN market has experienced tremendous growth. This is due to a confluence of factors including the adoption of industry standards and interoperability testing, the progressing of wireless LAN equipment to higher data rates, rapid decreases in product prices, and an industry shift toward mobility and the use of laptops. In 2002 the In-Stat market report [4] showed a strong shift of wireless LAN operation frequency from 2.4 to 5 GHz. The 5 GHz band offers the advantages of higher data rates, far more available spectrum, less sharing

with other uses such as cordless phones and Bluetooth radios, and, probably most important, an environment with much less noise and interference from other electronic devices.


This article first presents an overview of the 802.11a wireless LAN standard [1]. Following that, a transceiver architecture that employs two chips, one front-end analog chip and one baseband processing chip, that implement the 802.11a wireless LAN standard operating at 5 GHz is discussed. Also covered in this article are the architectural and design trade-offs made at the early stage of the design process that made possible this highly efficient solution in terms of both the design effort involved and the total build of material (BOM) cost.

BACKGROUND: 802.11A PHYSICAL LAYER

Wireless networking systems can be best understood by considering the physical (PHY) and media access control (MAC) layers separately. The PHY layer of 802.11a is based on orthogonal frequency-division multiplexing (OFDM) [2], a modulation technique that uses multiple carriers to mitigate the effects of multipath. OFDM distributes the data over a large number of carriers that are spaced apart at precise frequencies.

As indicated in Fig. 1, the 802.11a standard supports multiple 20 MHz channels, with each channel being an OFDM modulated signal consisting of 52 carriers. Among the 52 carriers, 48 carry data and four carry pilot signals used for tracking. Each carrier is 312.5 kHz wide, giving raw data rates from 125 kb/s to 1.125 Mb/s per carrier depending on the modulation type employed — binary phase shift keying (BPSK), quaternary PSK (QPSK), 16-quadrature amplitude modulation (QAM), or 64-QAM — and the amount of error-correcting code overhead (1/2 or 3/4 rate code). The composite signal therefore has a data rate of 6 Mb/s up to 54 Mb/s in a 20 MHz channel.

OFDM is one of the most spectrally efficient data modulation techniques available. Instead of separating each of the 52 carriers with a guard band, OFDM overlaps them. If done incorrectly, this could lead to an effect known as intercarrier interference, where the data from one carrier



■ **Figure 1.** Channel allocation of the IEEE 802.11a standard within the 5 GHz band.

cannot be distinguished unambiguously from its adjacent carriers. OFDM avoids this problem by making sure that the carriers are orthogonal to each other by precisely controlling their relative frequencies and timing. This high degree of precision requires the receiver design to be highly robust to both frequency offset and timing offset, as discussed in later sections.

BACKGROUND: 802.11 MAC LAYER

Access methods for wireless channels fall into three general categories: contention methods, polling methods, and TDMA methods. 802.11 is based primarily on contention methods, with some polling capabilities as well. Contention systems such as 802.11 use heuristics, such as random backoff, listen-before-talk, and mandated interframe delay periods, to avoid (but not completely eliminate) collisions among individual users on the shared wireless medium. The 802.11 MAC also employs a beacon message that can be asserted by an access point (AP) to individually poll selected stations for sending or receiving data at specific times. The duration of the polling period is controlled by a parameter set by the AP and contained within the beacon message.

Another MAC layer consideration is whether there is a dedicated central controller such as an AP or base station. 802.11 uses an AP, but has a fallback method for when there is no centralized controller (the ad hoc mode). However, the operation of the network is more efficient with an AP present.

DESIGN CHALLENGES

In this section some of the early design decisions are reviewed, which eventually would determine not only the overall system performance, but also the amount of design effort required and the total system cost.

OFDM PEAK-TO-AVERAGE RATIO AND PA EFFICIENCY

OFDM modulation, which is highly desirable because of its resilience to multipath interference, can substantially complicate the transceiver design. Since the 802.11a OFDM symbol consists

of 52 carriers, its amplitude peak-to-average ratio can be as high as 17 dB ($10 \cdot \log(52) \sim 17$ dB). However, such extreme peaks are in practice rare and brief. It is not necessary to preserve these extreme peaks in order to demodulate the signal correctly. Therefore, both the transmitting and receiving circuitry are allowed to clip or compress the signal during its extremes in practical OFDM transceiver design.

A large peak-to-average ratio is most expensive to support in terms of cost and power in the power amplifier (PA). There are three limits that must be considered when choosing the appropriate power backoff from a PA's saturation power output (P_{sat}): the 802.11a transmit spectral mask, the transmit signal accuracy (error vector magnitude, EVM), and the regulatory out-of-band transmission limits. The 802.11a transmit spectral mask and the required EVM limits can be found in [1]. At lower data rates (6–24 Mb/s), the EVM required by 802.11a is less stringent, so the spectral mask is the limiting factor. At data rates above 24 Mb/s, the EVM requirement becomes the limiting factor (no greater than -25 dB is required at 54 Mb/s), so larger backoffs are necessary. Such large backoffs can easily result in poor PA efficiency. Two techniques that reduce the effective peak-to-average ratio of OFDM symbols are implemented in the digital baseband processing chip, as will be described later.

DYNAMIC RANGE REQUIREMENTS

The spectral efficiency of the 802.11a standard comes at the expense of a much larger dynamic range requirement on both the transmitter and receiver design. For example, the use of 64-QAM modulation requires a signal-to-noise ratio (SNR) of approximately 25 dB to achieve a bit error rate of 10^{-5} , which is substantially greater than that required by the frequency shift keying (FSK) modulation in Bluetooth and the QPSK modulation in 802.11b. This high SNR translates to stringent phase noise requirements for the frequency synthesizer and tight I/Q matching constraints for both the transmitter and receiver. This higher SNR also dictates the amount of quantization error allowable in the digital detection circuit.

Access methods for wireless channels fall into three general categories: contention methods, polling methods, and TDMA methods. 802.11 is based primarily on contention methods, with some polling capabilities as well.

Since the intended host of our design is a PC or a laptop, there exists a host CPU to perform non-timing critical functions such as association and authentication exchanges or data frame preparation.

Two types of receive dynamic range requirements are considered. The first is the dynamic range of the desired signal due to the receiver being close to or far from the transmitter. This type of dynamic range requirement can be satisfied using variable gain amplifiers in an automatic gain control (AGC) loop. 802.11a calls for a maximum receive signal size of -30 dBm and a minimum sensitivity of -82 dBm. However, -20 dBm to -91 dBm (71 dB desired signal dynamic range) is a more appropriate signal range for quality equipment. This wide range of variable gain requirement will be addressed later.

The second dynamic range requirement is the instantaneous dynamic range of the received signal itself. The 802.11a specification calls for blocking signals that can be 16 dB stronger than the desired in-band signal in the adjacent channel and 32 dB stronger in the alternate channel (for 6 Mb/s frames). These blocking signals cannot be allowed to saturate the receive chain, or the desired signal will become corrupted.

In addition to not saturating the receive chain by blocking signals, the received signal must be sized at each stage to ensure that sufficient SNR is preserved. The 54 Mb/s data frame is the most demanding. Theoretically this signal can be received with as little as 20 dB SNR with the coding gain afforded by error correcting codes. But 30 dB is a more appropriate practice to ensure that circuit noise will not be a limiting factor in detection performance. In addition, it is desirable to ensure that the noise of any stage in the receive chain is not the limiting factor. The dynamic range required in the radio frequency (RF) and intermediate frequency (IF) before the analog channel select filters remove part of the blocking signals can be calculated as

- 30 dB (SNR)
- + 15 dB (alternate blocker for 54 Mb/s)
- + 10 dB (peak-to-average power backoff)
- + 4 dB (implementation loss) = 59 dB

59 dB is an extremely high instantaneous dynamic range for some of the stages in the receive chain. It is therefore important to have the signal optimally sized at each stage. Our design employs adjustable gain in nearly every stage, and controls the gain of the stages intelligently from the digital baseband processor, as described later.

DIGITAL COMPENSATION OF ANALOG IMPAIRMENTS

Advances in complementary metal oxide semiconductor (CMOS) technology created an opportunity for integrated 5 GHz RF design. One main goal of our implementation of the 802.11a wireless LAN standard is to recast the traditional RF transceiver design principles based on discrete components to an architecture that favors integration. However, CMOS technology was developed for digital switching circuits, in which circuit noise was not a major target for optimization. Integrating high-accuracy RF components in CMOS therefore requires some modifications, or joint design, of both the analog and digital circuitry.

Digital CMOS technology provides fast switching capabilities that favor wideband

transceiver architectures. 802.11a uses a channel of approximately 20 MHz, and therefore is a natural candidate for CMOS integration. In the baseband design, due to adjacent channel blockers at the receiver and a very stringent spectral mask requirement at the transmitter, the sampling rate of the analog-to-digital converters (ADCs) in our design is set at 80 MHz, four times the received signal bandwidth, and the sampling rate of the digital-to-analog converters (DACs) is set at 160 MHz, eight times the transmitted signal bandwidth. Also, due to the requirement for a fairly large signal dynamic range, these ADCs and DACs need to be of 9-bit resolution, with measured peak SNR of 48 dB and 52 dB, respectively. These design specifications are made because accurate analog filters are difficult and expensive to integrate on chip, while CMOS ADCs/DACs can handle high-resolution and wideband sampled data at minimal power and cost. Consequently, most of the adjacent channel blockers and out-of-band transmitted signals are filtered out using digital filters, which can be better controlled and consume little power.

In RF transceiver design, it was understood that CMOS technology might introduce more circuit noise into the transmit/receive chains than that of discrete components. It was also desired that the analog RF transceiver chip have an extremely high yield despite the process and temperature variations that tend to degrade the performance of analog circuits. Extreme care was taken in the design and layout of the analog RF circuitry to allow for a wide margin to accommodate different process corners of the standard digital CMOS technology.

Various noise sources such as $1/f$ noise and phase noise are modeled and verified by measurement. These models are used to drive the digital algorithm design in an all-encompassing simulation environment. Wherever possible, sophisticated digital signal processing techniques are used to monitor, track, and compensate for the analog impairments that cannot be designed out during the circuit design phase. Examples of these analog impairments that rely at least partially on digital compensation are carrier frequency offset, sampling frequency offset, timing offset, phase noise, DC leakage current, and I/Q mismatch. The total silicon area used for correcting analog impairments is estimated to be 3 mm², with power consumption of approximately 25 mW.

MAC IMPLEMENTATION

The MAC layer functions have traditionally been implemented using a programmable processor such as an ARM processor. This approach has two major drawbacks. First, a programmable processor consumes a disproportionate amount of power compared to a dedicated hardware design, usually on the order of tenfolds. Second, to implement the 802.11a MAC protocol, a programmable processor would need to act within microseconds of an event or at precise intervals. This degree of timing accuracy casts a burden on the programming task. As a result, most of the MAC programs are written in assembly code, which requires an enormous amount of time and

effort, not to mention the risk incurred in the production schedule.

Since the intended host of our design is a PC or laptop, there exists a host CPU to perform non-timing-critical functions such as association and authentication exchanges or data frame preparation. A natural architecture to implement the 802.11 MAC would be to design a dedicated state machine for timing-critical functions and leave the non-timing-critical functions to be performed on the host CPU.

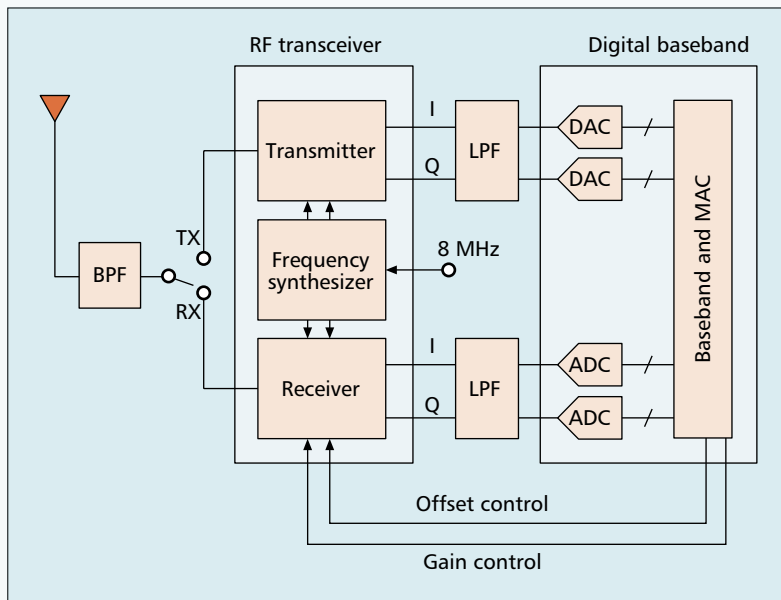
This hardware MAC concept is easier said than done, as all conventional wisdom indicates that programmability is the key to success because of its seeming flexibility. Eventually a hardware MAC architecture was chosen that proved to be the most power-efficient solution of all the MAC implementations ever reported. For example, a popular processor choice for performing the 802.11a/b/g MAC functions is the ARM922T implemented in 0.18 μm running at 200 MHz. This processor consumes 160 mW, not including the power for external memory and interface, while our hardware MAC consumes only 24 mW in 0.25 μm technology. Interestingly enough, the design time of this hardware MAC turned out to be relatively short compared to that of a processor/software approach.

SYSTEM PARTITION

Since the decision to design two chips, rather than integrating all functions into a single chip, was made early in the design process, an appropriate interface between the two chips must be determined. A digital interface would allow the ADCs/DACs to reside on the analog chip, simplifying the testing and portability issue of the digital baseband processing chip. However, a digital interface would require at least 36 more pins for communication between the two chips, as both transmit and receive chains generate in-phase and quadrature signals. Furthermore, the relatively large switching power in the ADCs/DACs might introduce sufficient substrate noise to disturb the receiver front-end operation with input signals as small as a few microvolts. Therefore an analog interface is used, which minimizes the pin count and delivers low-noise high-resolution analog signals between the analog RF transceiver chip and the baseband processing chip.

Figure 2 shows the system partition of the two-chip set implemented in 0.25 μm CMOS. The analog baseband I and Q input signals for the RF transmitter are generated by two 160 MHz 9-bit current-steering DACs on the baseband chip. On the receiver side, the quadrature signals I and Q at the RF transceiver output are digitized by two 80 MHz 9-bit ADCs on the baseband chip before being processed by the baseband and MAC processor. Off-chip LC low-pass filters are used between the two chips for noise bandwidth limitation and anti-aliasing reasons. No off-chip intermediate frequency (IF) filter is required in the system.

The next two sections of the article describe the architecture and implementation of the analog RF transceiver [5] and the baseband processing chip [6], respectively. The circuit implementation of individual blocks and experimental results are also reported.



■ Figure 2. IEEE 802.11a wireless LAN system architecture.

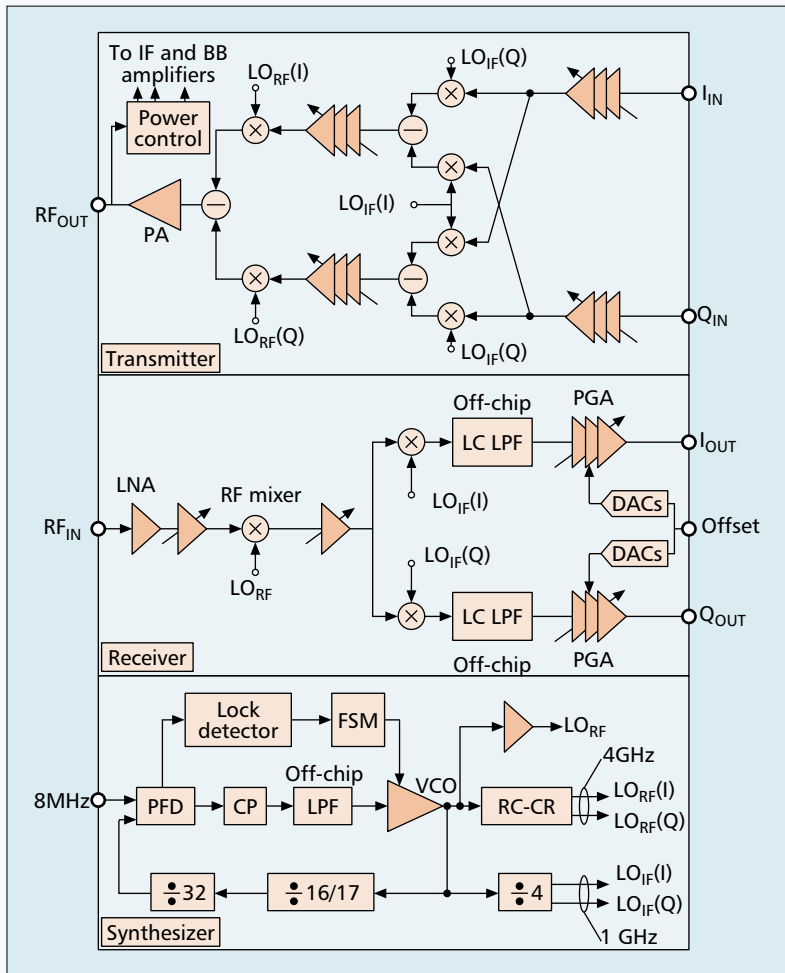
ANALOG RF TRANSCEIVER DESIGN

The analog RF transceiver consists of a transmitter, a receiver, and a frequency synthesizer. We will describe the architecture and frequency plan of the RF transceiver design first.

ARCHITECTURE AND FREQUENCY PLANNING

The architecture and frequency plan of the RF transceiver play an important role in the complexity and performance of the overall system. Two of the most common choices in transceiver architecture are direct conversion and the traditional superheterodyne. Direct conversion is usually preferred in a fully integrated design because it avoids the need for an off-chip IF filter and requires only a single frequency synthesizer. However, it suffers from drawbacks such as local oscillator (LO) leakage and frequency pulling due to the fact that the synthesizer operates at the same frequency as the RF signal. The superheterodyne architecture overcomes many of the disadvantages of direct conversion at the expense of an off-chip IF filter and an extra frequency synthesizer.

The RF transceiver described in this article uses a dual conversion architecture with a sliding IF of 1 GHz [5]. Shown in Fig. 3 is the detailed block diagram of the transceiver. On the transmit side, the baseband I and Q signals are first mixed to 1 GHz by a pair of image-reject mixers. Each mixer is a doubly balanced configuration and generates a single upper-sideband signal in each of the I and Q mixer outputs. The undesired lower-sideband signal is attenuated by over 40 dB [7]. The quadrature 1 GHz IF signal is then converted to 5 GHz by the RF mixer. Double image reject mixers are used in the transmitter in order to avoid the need for an IF filter. The upconverted 5 GHz signal is finally transferred to the antenna through an on-chip PA. Since the transmitter output signal is not at the frequency of the LO, no LO pulling is caused by the transmitter PA. Furthermore, any LO leakage to the antenna



■ **Figure 3.** RF transceiver chip block diagram.

will be at least 1 GHz away from the in-band signal and appear as out-of-band tone, and will not interfere with the operation of other receivers operating in the 5 GHz band.

The receiver frequency plan is very similar to that of the transmitter. An incoming 5 GHz RF signal is first mixed down to IF at 1 GHz and then converted to the baseband quadrature signals. For an RF signal centered at 5 GHz, the image channel is located 2 GHz away at 3 GHz. This undesired signal will be attenuated at least 23 dB by the bandpass gain stages between the receiver input and RF mixer. By mixing the incoming RF signal with the 4 GHz LO, the RF mixer generates the desired 1 GHz IF and a spurious signal at 9 GHz. The spurious signal is attenuated by the inherent bandwidth limitation of the circuits. As a result, an image reject mixer is not required in the receiver.

The use of a sliding IF architecture, whereby the LO_{IF} is generated from the LO_{RF} using a divide-by-four counter, eliminates the need for two synthesizers. Designed in a twisted ring architecture [8], the divide-by-four counter can inherently provide very precise quadrature LO signals at 1 GHz, thereby improving the transmitter's image rejection.

Advantages and challenges accompany the implementation of the RF transceiver in CMOS technology. CMOS ultimately provides a signifi-

cant cost advantage over alternatives. Moreover, scaled CMOS processes generally offer multiple layers of interconnects that allow the use of integrated inductors with quality factors as high as 10 at a frequency of 5 GHz [9]. These inductors are used extensively throughout the transceiver described in this work in order to enhance the gain of narrowband amplifiers. Device characteristics in CMOS can vary significantly over changes in process and temperature, resulting in substantial variations in performance. This drawback can be overcome by using adjustable gain stages and implementing an automatic gain control (AGC) algorithm in the digital domain.

TRANSMITTER

The design of the power amplifier (PA) is one of the most challenging tasks in transmitter implementation. The large peak-to-average ratio of OFDM symbols requires the PA to provide substantially higher peak output power than its average. In practical applications, however, since extreme signal peaks are infrequent, the peak-to-average ratio requirement becomes a function of the modulation used and the degree of non-linearity in the PA power transfer curve.

Our design uses a three-stage class A PA in the transmitter, wherein each stage consists of a cascaded differential pair. The gate terminals of the cascode transistors are biased at the supply voltage. The cascode topology allows the PA to use a 3.3 V supply for increased headroom and improved linearity. Wherever possible on-chip inductors are used to form parallel resonances with gate or parasitic capacitance to increase gain and noise immunity. The fully differential PA output reduces the effects of parasitic supply and ground inductances. Closed loop power control provides a constant transmitted output power independent of process, temperature, and supply voltage variations. The power control loop, consisting of a peak detector, a comparator, and 24 dB of adjustable transmitter gain in 0.5 dB steps, adjusts the transmitter gain until the PA output matches a preprogrammed level. The transistors are sized to deliver a measured saturated output power (P_{sat}) of 22 dBm.

Measurements indicate that the transmitter can provide a peak output power of 22 dBm and an average OFDM output power of 17.8 dBm. The measured transmit output power is a function of the transmit signal data rate. At low data rates (6–24 Mb/s), the measured output power is about 18 dBm and is limited by the spectral mask requirements. At higher data rates, an additional backoff of 6 dB is needed for the transmitter to meet the -25 dBc EVM requirement for transmitting 64-QAM OFDM signals.

RECEIVER

The receiver LNA consists of a cascaded differential pair with inductive loads that tune the amplifier output to 5 GHz. The inductive degeneration formed by on-chip inductors results in a complex input impedance that can be matched to a 50- Ω source impedance with an off-chip matching network.

The receiver mixes the 5 GHz RF input first to the 1 GHz IF and then to the quadrature baseband outputs for digitization by the ADCs

on the baseband chip. The entire receive chain is designed to provide sufficient dynamic range and linearity for receiving 64-QAM OFDM signals, which requires at least 30 dB of SNR. The RF and IF gain stages have a maximum combined gain of 36 dB that significantly reduces the noise contribution of subsequent baseband stages.

The downconverted I and Q signals are passed through the off-chip passive LC channel-select filters and amplified by a programmable gain amplifier (PGA). The off-chip channel-select filters help to reduce the adjacent channel blockers by approximately 6 dB and therefore reduce the dynamic range requirements of the following ADCs. The PGA comprises a cascade of three stages with a composite gain that can be adjusted from 0 to 41 dB in 1 dB steps. The DC offset of the receive chain is cancelled using two pairs of 6-bit DACs. The DC offset cancellation, AGC, frequency offset cancellation, timing offset correction, and receive signal strength indicator are all implemented with digital algorithms in the baseband chip, as described in the next section.

The receive chain achieves over 75 dB of adjustable gain. The I and Q baseband signals measured at the receiver output indicate an I/Q phase mismatch of 1.5° and amplitude mismatch of 1.5 dB without the use of any calibration. The measured noise figure of the entire receive chain from LNA to baseband PGA is 8 dB.

SYNTHESIZER

The frequency synthesizer generates the quadrature 1 GHz and 4 GHz LO frequencies needed for the mixers in the receive and transmit chains. The synthesizer phase locks an on-chip VCO to an 8 MHz reference. The VCO frequency is fine tuned using two P+/N-well varactors. The switching sequence is determined by a state machine in conjunction with a lock detector circuit. For a particular RF channel, if the varactors fail to force the loop to lock, the state machine switches in enough fixed capacitors until the varactors can pull the loop to lock.

The variable divider in the feedback loop consists of a divide by 16/17 dual-modulus prescaler followed by a divide by 32 and a channel select decoder. The synthesized frequency can be varied from 4.128 to 4.272 GHz, which corresponds to an RF carrier frequency ranging from 5.16 to 5.34 GHz. The quadrature 1 GHz LO signals are generated by a divide by four counter. Designed in a twisted ring architecture, this divider can generate very precise 1 GHz quadrature signals and maintain this precision over process and temperature. The quadrature 4 GHz LOs are generated by passing the VCO waveform through a single-stage RC-CR polyphase filter [10].

The composite phase noise of the synthesizer measured at the output of the power amplifier shows that the close-in phase noise is -87 dBc/Hz at 1 kHz frequency offset. Outside the loop bandwidth of 250 kHz, the phase noise decreases to -112 dBc/Hz at 1 MHz offset.

POWER CONSUMPTION

The analog RF transceiver has been integrated in a 0.25 μm , single-poly, 5-metal CMOS technology. It occupies a total area of 22 mm² and is

packaged in a 64-pin leadless plastic chip carrier with an exposed backside contact for good thermal and electrical performance. A die photo of the chip is shown in Fig. 6a. The transceiver operates from a 2.5 V supply with 3.3 V I/O. The transmit chain dissipates 790 mW of power including a 22 dBm power amplifier. The receiver and synthesizer consume 250 mW and 180 mW, respectively.

BASEBAND AND MAC PROCESSOR DESIGN

The design goals of the baseband and MAC processor chip are to achieve close to theoretical physical layer performance, apply signal processing algorithms to compensating for analog impairments, and deliver an extremely power-efficient design.

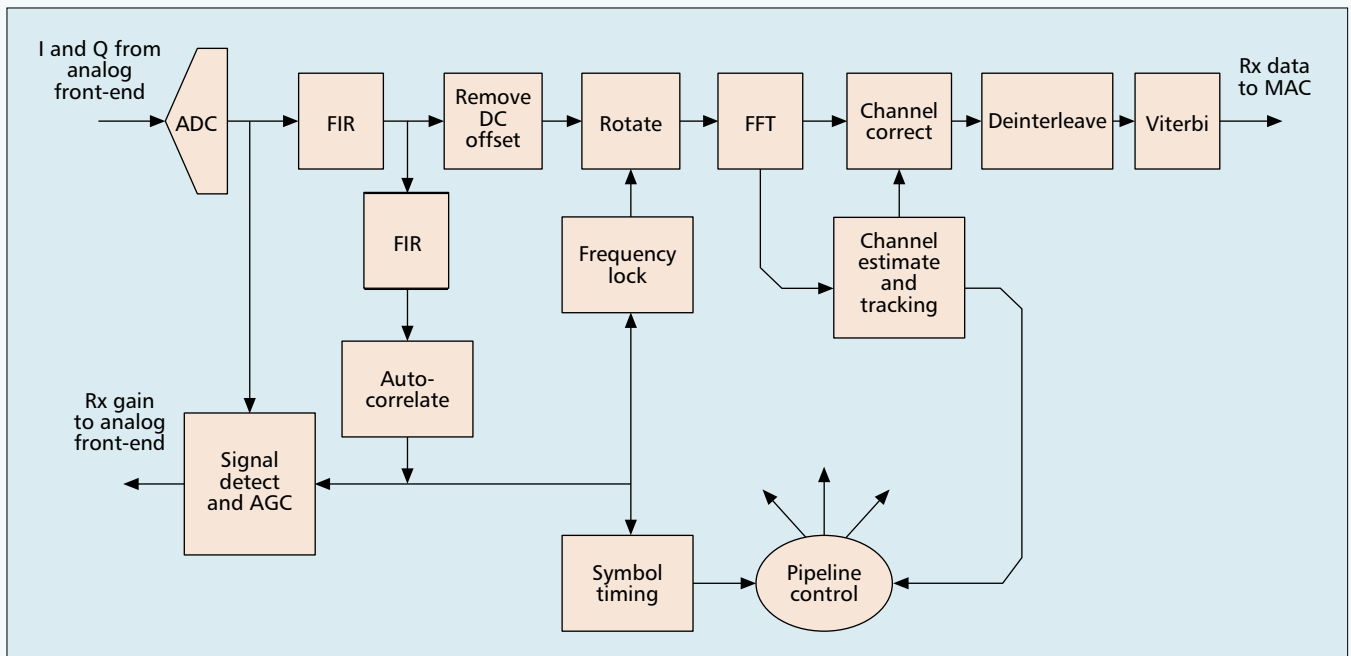
The baseband and MAC processor contains three units: the baseband transceiver that implements the 802.11a physical layer function, the protocol control unit (PCU) that manages all low-level timing-critical aspects of the 802.11 MAC layer function, and the host interface unit (HIU) that provides connectivity to the host processor over a PCI bus. A DMA engine manages the transfer of frame data and control information between the PCU and HIU. In this article we focus on the design of the baseband transceiver and PCU.

BASEBAND TRANSCEIVER

The baseband transmitter generates OFDM symbols according to the 802.11a specification. A 128-point IFFT reduces the filter length of the following transmit interpolation filter, thereby preserving the guard interval for multipath. Dual 9-bit 160 MHz DACs employ a current steering structure and pass digital baseband data to the analog transceiver. The required resolution of the DACs was derived as follows. A 9-bit DAC has a dynamic range of 54 dB. When combined with oversampling, it gives a dynamic range of ~64 dB. In the digital domain we might back the transmit signal off by as much as 10 dB from full scale. Therefore, the quantization noise floor would be ~54 dB below the signal level. The 802.11a specification requires the transmit spectral mask to be 40 dB down in the alternate channel, and the FCC out-of-band emission requirements are even more stringent. In practice an 8-bit DAC would be sufficient given this much oversampling, but 9-bit resolution provides a good margin for other circuits to consume the noise budget.

To reduce the effective peak-to-average ratio of the transmitted OFDM symbol for a nonlinear PA, two digital predistortion techniques are implemented. The first employs programmable scaling after the interpolation filter that trades DAC quantization noise against increased probability of clipping. For low-rate data frames (6–24 Mb/s), OFDM symbols can be clipped substantially without violating the EVM requirement of the 802.11a specification. For high-rate data frames (36–54 Mb/s), which require much higher peak-to-average requirement for reliable detection, a second predistortion technique is

The design goals of the baseband and MAC processor chip are to achieve close-to-theoretical physical layer performance, apply signal processing algorithms to compensating for analog impairments, and delivers an extremely power-efficient design.



■ **Figure 4.** Receiver architecture for baseband processing.

introduced. This technique uses a lookup table to dynamically scale up samples with large amplitude, which are expected to be compressed by the nonlinearity of the PA to create an effectively linear output signal.

The baseband receiver, shown in Fig. 4, contains dual 9-bit 80 MHz ADCs that cover an input range of ± 500 mV. The oversampling relative to the channel bandwidth of 20 MHz simplifies the anti-aliasing filter design and allows the filtering of adjacent channel blockers to be done primarily in the digital domain.

Calibration of DC offset and gain in the analog receiver is performed using digital algorithms. The AGC maximizes received signal size at the ADC input while providing headroom for adjacent channel interference and the peak-to-average ratio of received OFDM symbols. The relatively short time for AGC (approximately 4 μ s) in 802.11a demands a quick loop from digital power measurement to analog gain adjustment. The receiver gain under AGC control is composed of RF, IF, and baseband PGA stages, as well as an antenna switch that can be opened during receive to provide additional attenuation in the event of a very large RF input signal.

The ADC outputs pass through lowpass downsampling filters that eliminate possible adjacent channel signals. Signal detection, frequency offset estimation, and symbol timing all rely on auto-correlation of the periodic training symbols provided in the preamble. The short preamble symbols, 10 periods of 0.8 μ s each, are used to detect the existence of a frame, calculate the carrier frequency offset that is fed into the frequency rotator, and estimate the symbol timing. The long preamble symbols, two periods of a training OFDM symbol each 4 μ s long, are averaged, fast Fourier transformed (FFTed), filtered, and inverted to form an inverse of the channel estimate in the frequency domain. A 128-point FFT reduces filtering requirements for

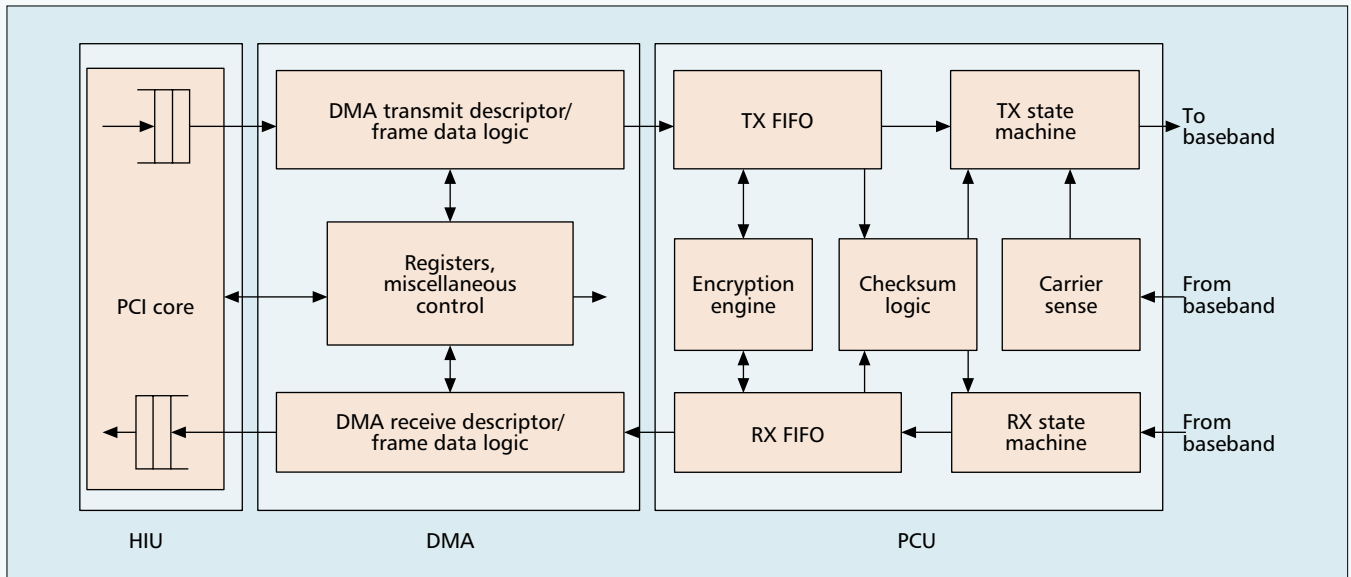
eliminating adjacent channel signals, preserves the guard interval, and shares hardware with the inverse FFT (IFFT) used in the transmitter.

For each of the four pilots in a data symbol, the phase with respect to that of the training symbol is computed. A least squares fit of the four pilot phase differences determines the phase adjustment for each data carrier. Pilot phase monitoring can track frequency offset estimation error, phase noise, and symbol timing drift. Pilot magnitude tracking compares pilot power in the data symbols to the training symbols and monitors gain variations. Equalized data is passed to a Viterbi decoder using a radix-4 fully parallel soft decision traceback architecture.

MAC ARCHITECTURE

The MAC architecture consists of the PCU, a DMA engine, and the HIU (Fig. 5). The PCU manages all low-level timing-critical aspects of the 802.11 MAC specification. It formats and sends outgoing data frames to the baseband transmitter and processes incoming data frames from the baseband receiver.

Timing-critical functions require the MAC to act within microseconds of an event or at precise intervals. Examples of these functions include the channel access mechanism such as random backoff and listen-before-talk, channel state and networkwide timer updates, checksum generation and verification, hardware-level frame retry, and generation of special frames such as periodic beacons and receive acknowledgments. Non-timing-critical functions are performed in the driver software executing on the host CPU. These include complex frame exchange sequences (e.g., association and authentication exchanges), frame fragmentation and defragmentation, frame buffering and bridging, and other network management portions of the 802.11 protocol.



■ **Figure 5.** MAC architecture.

The driver software builds transmit frames as a collection of frame fragments in host memory and passes a pointer to a corresponding list of transmit descriptors to a DMA engine. The DMA engine traverses the list and performs the necessary data fetches from host memory, passing the coalesced frames to the PCU. The DMA engine provides full scatter/gather capability, including support for arbitrary byte alignment and byte lengths, to avoid multiple data copy operations on the host.

The PCU implements multiple encryption methods, including the traditional Wireless Encryption Protocol (WEP) based on RC4. The PCU encrypts the frame if encryption is enabled and generates the proper checksum value. It follows the 802.11 carrier sense multiple access with collision avoidance (CSMA/CA) access procedure to gain access to the channel and then forwards the frame to the baseband transmitter. During receive, the PCU extracts the frame type, verifies the checksum, and generates a response frame (typically a receive acknowledgment) if appropriate. The PCU passes the received frame to the DMA engine, which interprets a series of descriptors to transfer the frame data to host memory.

The PCU also provides power management functions. Acting as a station in a multinode network, for example, the wireless LAN chip set can be programmed to sleep automatically and awake just before the next beacon is scheduled to arrive. The PCU parses the incoming beacon to determine whether to remain awake for additional frames or resume sleep.

The MAC architecture is implemented using dedicated control and data path logic, and includes registers that allow host software to configure and control its operation. This yields an overall design that is compact, power-efficient, and requires no off-chip RAM or program storage, yet is flexible enough to accommodate the vagaries of the 802.11a protocol as well as the additional needs of the host operating system and driver software.

The baseband processing chip occupies a sili-

con area of 46 mm² and consumes 326 mW during active transmit and 452 mW during active receive. Power is reduced by utilizing 73 gated clock trees with independent enables.

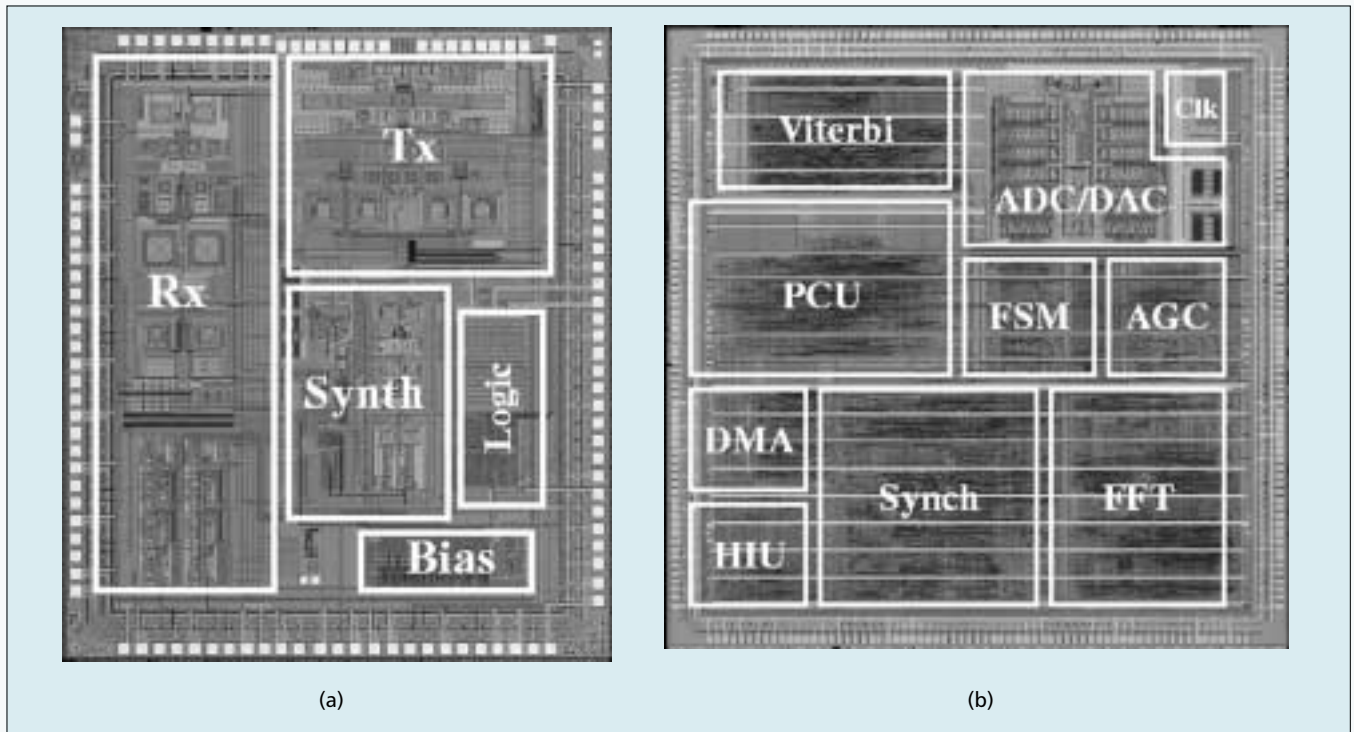
Along with the companion 5 GHz RF transceiver chip, the baseband and MAC processor chip is fully compliant with the 802.11a standard and far exceeds all mandated performance requirements. One example is the much lower (5–10 dB lower) receiver sensitivity achieved for all data rates. Additional data rates up to 108 Mb/s are supported by varying the internal clock frequencies and adjusting transmit and receive filters. The oversampled ADC and DAC designs accommodate the higher data rates. Figure 6 shows the die photo of the two-chip set.

CONCLUSIONS

In this article we discuss various design trade-offs and circuit implementation of the 802.11a wireless LAN standard in a two-chip set using 0.25 μm CMOS technology. Because of this design, wireless LAN systems are now available at a much lower price, and deliver much higher data rates and performance. Compared to other wireless LAN standards such as 802.11b or 802.11g, 802.11a is becoming popular due to its high throughput, large system capacity, and relatively long range [11]. Additional spectrum is being allocated for wireless LAN use at 5 GHz. New modulation techniques such as that used for antenna beamforming will increase data rates beyond 100 Mb/s, and extend ranges to several kilometers. Enhancements to the 802.11 MAC will provide wireless quality of service, enabling high-quality wireless voice, video, and audio transmission. All of these improvements will create opportunities for wireless LAN beyond today's use in corporate and home data networks.

ACKNOWLEDGMENTS

The authors wish to acknowledge the whole engineering team at Atheros Communications for this accomplishment. In particular, we thank



■ **Figure 6.** a) Die micrograph of the analog RF transceiver chip; b) die micrograph of the baseband and MAC processor chip.

Rick Bahr for his leadership, Bob Brodersen for his insightful suggestions, and the engineers for their contributions to the design, layout, and testing of the chip set.

REFERENCES

- [1] IEEE Std. 802.11a-1999, "Wireless LAN MAC and PHY Specifications — High-Speed Physical Layer in the 5 GHz Band," ISO/IEC 8802-11:1999(E)/Amd 1:2000(E), New York: IEEE, 2000.
- [2] R. Van Nee and R. Prasad, *OFDM for Wireless Multimedia Communications*, Artech House, 2000.
- [3] IEEE Std. 802.11b-1999, "Wireless LAN MAC and PHY Specifications," ISO/IEC 8802-11:1999(E), New York: IEEE, 1999.
- [4] A. Noguee, "WLAN Chipset Market — The Incredible Journey Is Just Beginning," In-Stat rep. no. IN020271WY, Mar. 2002.
- [5] D. Su *et al.*, "A 5 GHz CMOS Transceiver for IEEE 802.11a Wireless LAN," *ISSCC Dig. Tech. Papers*, Feb. 2002, pp. 92–93.
- [6] J. Thomson *et al.*, "An Integrated 802.11a Baseband and MAC Processor," *ISSCC Dig. Tech. Papers*, Feb. 2002.
- [7] T. Lee, *The Design of CMOS Radio-Frequency Integrated Circuits*, New York: Cambridge Univ. Press, 1998.
- [8] E. J. McCluskey, *Logic Design Principles*, Prentice-Hall, 1986.
- [9] C. Patrick Yue and S. Simon Wong, "On-Chip Spiral Inductors with Patterned Ground Shields for Si-Based RF ICs," *IEEE J. Solid-State Circuits*, vol. SC-33, no. 5, May, 1998, pp. 743–52.
- [10] F. Behbahani *et al.*, "CMOS Mixers and Polyphase Filters for Large Image Rejection," *IEEE J. Solid-State Circuits*, vol. SC-36, no. 6, June 2001, pp. 873–87.
- [11] http://www.atheros.com/pt/atheros_range_whitepaper.pdf, Wireless LAN Range Performance.

BIOGRAPHIES

TERESA H. MENG [F] (meng@mojave.stanford.edu) received her Ph.D. in electrical engineering and computer science from the University of California (UC) Berkeley in 1988.

She joined the faculty of the Electrical Engineering Department at Stanford University in 1988, and was appointed the Reid Weaver Dennis Professor in 2003. In 1998 she took leave from Stanford and founded Atheros Communications Inc. Awards and honors for her research work at Stanford include an NSF Presidential Young Investigator Award, an ONR Young Investigator Award, an IBM Faculty Development Award, a best paper award from the IEEE Signal Processing Society, and the Eli Jury Award from UC Berkeley. She is the author of one book, numerous book chapters, and over 200 technical articles in journals and conferences.

BILL MCFARLAND joined Atheros Communications in 1999, and is currently director of algorithms. He manages a team developing digital signal processing algorithms, defines digital and analog radio architectures, and represents Atheros in regulatory and standardization efforts. Prior to joining Atheros, he spent 14 years at the Hewlett Packard Research Laboratory, working on high-speed digital test equipment and fiber optic communications links, and managed the wireless circuits research group. He has published over 25 papers, and holds eight patents.

DAVID SU is director of analog design at Atheros Communications, engaged in the design of CMOS transceivers for wireless LANs. Prior to joining Atheros in February 1999, he spent 10 years at Hewlett-Packard Company (IC Business Division and HP Labs) designing CMOS mixed-signal, analog, and RF ICs. He is also a consulting associate professor at Stanford University. He holds a Ph.D.E.E. from Stanford University as well as M.E. and B.S.E.E. from the University of Tennessee, Knoxville.

JOHN THOMSON received an M.S.E.E from UC Berkeley, and a B.S.E.E from the University of Alberta. He is a digital design manager at Atheros Communications, responsible for baseband design and verification, and physical design. Prior to Atheros, he designed high-performance pipelines at Chromatic Research and Sun Microsystems. His awards include the University of California Regent's Fellowship, Sir James Loughheed Award of Distinction, and the Henry Birks and Sons, Limited, Medal. He has co-authored several patents.