



ATHEROS®

**Achieving Higher Throughput and Greater Range
in 802.11n Networks by Sustaining Signals
for Improved Performance and Reliability**

ABSTRACT

This whitepaper introduces engineers and managers to the key technical issues associated with building reliable, high-performance 802.11n equipment. The Multiple Input Multiple Output (MIMO) approach to 802.11n introduces many design challenges to sustaining wireless signals, including limiting coding losses, achieving better synchronization between multiple receivers, and improving demodulation between MIMO signals. By implementing innovative Low Density Parity Check (LDPC) codes, Transmit Beamforming (TxBF), and Maximum Likelihood (ML) technologies, 802.11n SoC devices provide the performance and reliability required to drive efficient wireless links with the highest throughput at range.

Multiple Input Multiple Output (MIMO) technology forms the foundation of high-performance for 802.11n. Through the use of multiple transmit and multiple receive antennas, it becomes possible to apply a variety of techniques – including spatial multiplexing, beamforming, and advanced demodulation – to sustain signals over greater distances at higher data rates with increased reliability. In addition to transmitting multiple data streams simultaneously to increase overall throughput, systems need to be able to intelligently and dynamically concentrate the transmitted energy in the direction of the receiver as well as continuously adapt and operate robustly in environments with many reflections.

The 802.11n standard offers beamforming techniques to create a virtual array of antennas that form high-gain beams focused at clients. Focusing beams has the impact of increasing range and lowering overall environmental interference. Specifically, rather than blast energy in all directions, transmit energy is focused directly at the intended receiver. Such a focused beam ensures that the majority of the energy transmitted will reach the proper receiver. In addition, focused beams reduce the amount of energy sent in other directions and thereby cause less interference with other wireless links.

Maximizing the accuracy of these beams is crucial to achieving the highest effective throughput. Unless the signal can be sufficiently sustained from transmitter to receiver, errors and noise will erode throughput. For example, coding losses, lack of phase alignment, and marginal demodulation between multiple receivers results in higher bit error rates (BER) that lead to more retransmissions, wasted signal energy, and greater interference.

From a customer perspective, effective throughput and range are directly correlated to how well a signal has been sustained. Figure 1 shows how the combination of Low Density Parity Check (LDPC) codes, Transmit Beamforming (TxBF), and Maximum Likelihood (ML) equalization provides superior performance compared to systems which do not sustain signals using these technologies. The gains realized are significant and provide impressive market differentiation that even the most non-technical users can appreciate: either 50% greater range at ~25Mbps throughput or 66% higher effective bandwidth.

TCP Throughput (w/ 11% collision), Channel D, HT40, 32 A-MPDUs

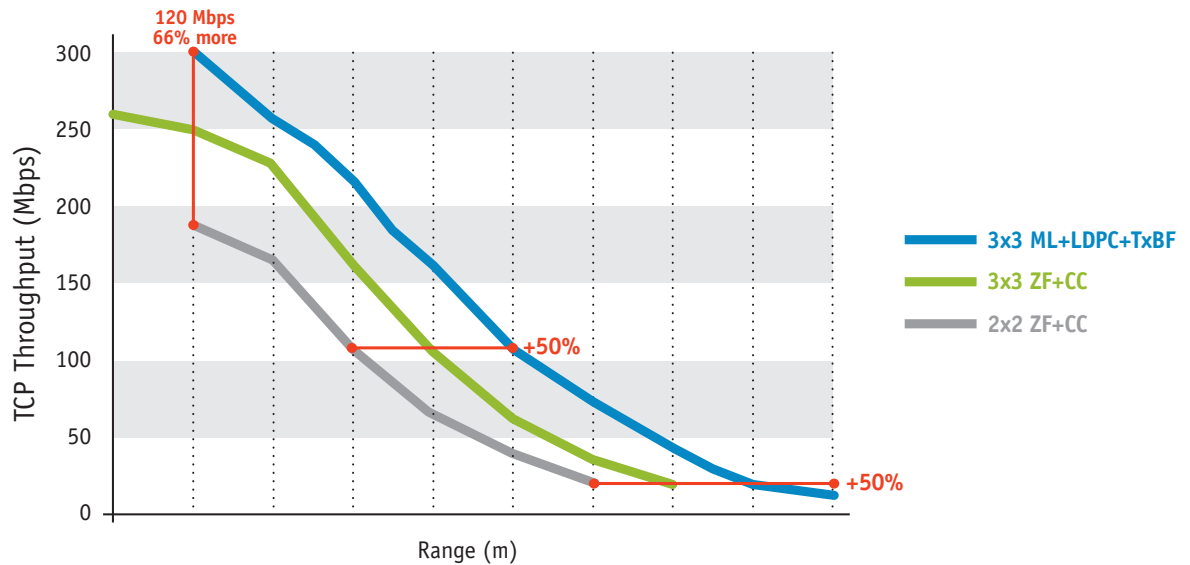


Figure 1: Effective Performance Improvement

Combining Low Density Parity Check (LDPC), Transmit Beamforming (TxBF), Maximum Likelihood (ML), and other technologies provides superior performance that even the most non-technical users can appreciate: either 50% greater range or 66% higher effective bandwidth, depending upon the particular operating conditions.

ELIMINATING CHANNEL LOSSES THROUGH LDPC

By its nature, wireless communication is vulnerable to a wide variety of environmental factors that inhibit the accurate transfer of data. Sources of errors include collisions with other Wi-Fi networks, interference from other appliances sharing the same spectrum, and other types of noise that adversely affect signal integrity. Errors manifest as incorrect bits in received data. If these incorrect bits cannot be resolved, the receiver must alert the transmitter to resend the data, resulting in increased latency and loss of effective bandwidth.

Wi-Fi employs convolutional coding techniques to address errors introduced between the transmitter and receiver. Convolutional coding as a technology has been around for more than 20 years and is a proven method for increasing system reliability by reducing the bit error rate (BER) of transmission links. By design, convolutional coding not only detects errors, but also corrects a limited number of bit errors across a packet.

The increased reliability of convolutional coding comes at the cost of imposing overhead. With 8B/10B encoding, for example, every 8 bits to be transmitted is converted to a 10-bit code, representing 20% overhead. Put another way, the cost for this form of encoding is 20% of the available bandwidth. Efficiency of coding, then, is a measure of the improvement in signal gain to the overhead required to achieve it.

Low Density Parity Check (LDPC) technology provides improved performance gains compared to standard convolutional coding. LDPC's key attraction is that it uses a high performance error correction mechanism that approaches the theoretical Shannon limit for efficiency. LDPC has been recognized as a viable technology for many years, and has been recently reinvented to exploit existing technologies. The standard algorithm is known as a belief propagation algorithm that uses an iterative decoding process. The primary difficulty with LDPC is that the decoder is particularly compute intensive and has a large implementation footprint, leading to both slow convergence and higher silicon cost. As such, it is unrealistic to implement in 802.11n applications.

Rather than impose the high overhead and cost of a fully implemented standard belief propagation algorithm, designers have the option of implementing what is known as a layered belief propagation algorithm (LBP). While the LBP algorithm does not provide the accuracy of a fully implemented belief propagation algorithm, it does provide superior performance compared to convolutional coding methods while balancing cost and speed. Cost is controlled as an LBP implementation requires a small footprint compared to other algorithms and architectures, while allowing for appropriate hardware pipelining so that processing speed is increased without compromising performance. The LBP algorithm also increases performance by employing a dynamic scaling factor which is chosen based on channel estimates to enable faster convergence.

Another key characteristic of an LDPC implementation is iteration control. Theoretically, the more iterations of the algorithm that are made, the better the performance. However, throughput and latency constraints limit the number of overall iterations reasonably possible. The use of partial iterations enables a cost-effective implementation of LDPC that optimally utilizes the available decoding time for each code word to maximize convergence within given system constraints. The cost savings through reduced die size and performance improvements obtained from the LBP algorithm are impressive. Simulations show a ~2 dB increase in performance gain compared to convolutional coding, given the same overhead.

ACHIEVING BETTER PHASE ALIGNMENT THROUGH TXBF

Transmit Beamforming (TxBF) is an exciting technology that enhances the reliability and performance of beamformed links by allowing the transmitter to generate signals that can be optimally received. Beamforming uses sounding techniques to align the transmitter with the receiver. The transmitter sends a signal (“Where are you?”) and listens for a response from the receiver (“I’m right here”). By changing the characteristics of the transmission, the transmitter can hone in on the receiver’s location to tune the beam to be as narrow as possible.

With historical (open loop) beamforming techniques, the transmitter is effectively forced to estimate where the receiver is and whether a change improves or degrades signal reception. When signals and reflections arrive in phase, they add their energy to create the strongest possible signal. However, when they arrive out of phase, they begin to destructively interfere with each other, reducing the signal energy reaching the receiver. Ideally, transmitted signals and their reflections arrive phase-aligned at the receive antenna.

Closed loop TxBF techniques improve accuracy by enabling the receiver to provide direct feedback to the transmitter to maximize the phase alignment of signals and their reflections. Closed loop TxBF opens a channel to the transmitter that allows the receiver to provide specific data on how well it is receiving signals. In this way, the transmitter can more quickly and accurately assess the optimal beam to use. The final result of closed loop TxBF is an accurate steering matrix that is applied to signals before transmission.

Figure 2a

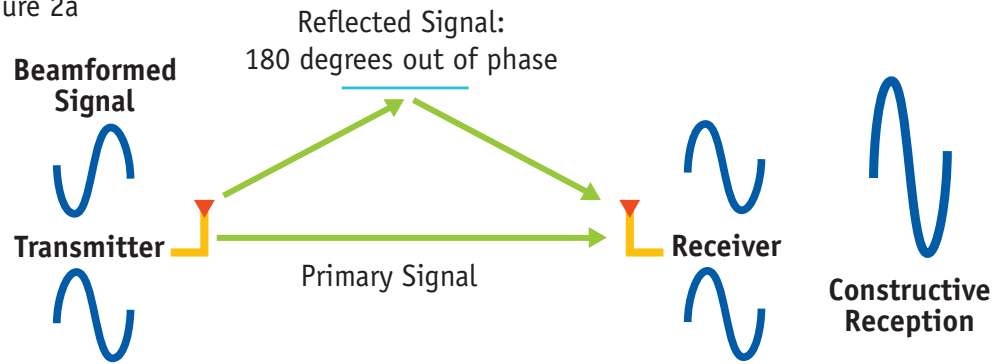
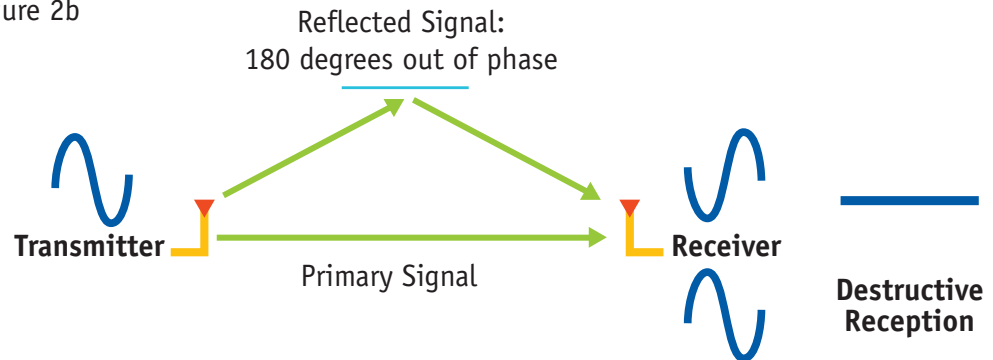


Figure 2b



Figures 2a and 2b

Accurate phase alignment of signals and their reflections is critical to maintaining the highest throughput. 2a) When signals and reflections are in phase, they constructively interfere with each other, adding their energy to improve signal reception. 2b) When signals and reflections are out of phase, destructive interference reduces the signal energy that reaches the receiver.

The 802.11n standard specifies four forms of TxBF: one implicit method and three explicit methods. With implicit TxBF, the steering matrix is created by the transmitter by tracking training symbols or High Throughput Long Training Fields (HT-LTF) and assuming the channel is reciprocal. Radio calibration is required for this method as the transmitter needs to calibrate the difference between its transmit and receive chains. This calibration is achieved using feedback from the receiver.

Explicit TxBF operates by having the receiver make a direct channel estimate from HT-LTFs sent by the transmitter and sending this feedback to the transmitter. The steering matrix can be calculated by either the transmitter or receiver. The 802.11n standard specifies three types of feedback:

- Channel state information: the receiver sends channel information to the transmitter to compute the steering matrix
- Non-compressed beamforming: the receiver calculates and sends the steering matrix to the transmitter
- Compressed beamforming: the receiver sends a compressed steering matrix to the transmitter

The advantages of the implicit method are that it requires only a light load on the receiver and has little over-the-air overhead. With the explicit method, no calibration is required. While an access point can use any of the four methods, because channel state information TxBF consumes a great deal of airtime with its overhead and brings down performance, few if any products are expected to use this method. As a result, 802.11n products will only need to support implicit TxBF as well as non-compressed and compressed explicit TxBF.

The 802.11n standard defines how and what data TxBF passes. How the information is processed can be implemented in a proprietary fashion.

MORE ACCURATE DEMODULATION THROUGH ML TECHNOLOGY

A similar performance/cost trade-off can be made when implementing equalizer technology. MIMO spatial multiplexing entails sending multiple data streams from the various transmit antennas at the same time. These streams arrive at the receiving antennas overlaid upon each other. A MIMO equalizer is used on the signals from the multiple receive antennas to separate the data streams into their original independent forms. Current implementations use the Zero Forcing (ZF) MIMO equalization method. New Maximum Likelihood algorithms (ML) offer greater accuracy.

As with LDPC, the standard ML algorithms are too computation-intensive to be cost-effectively implemented; the complexity of even a simple 2x2 system with 64QAM is prohibitive. To best achieve the best performance while controlling cost and managing complexity, spatial multiplexing needs to be implemented using a more cost-effective and balanced approach.

The advantage of a simplified maximum likelihood algorithm is that it achieves close to optimal performance but has a substantially lower implementation cost compared to a full ML implementation or other conventional suboptimal algorithms, including sphere decoding, which suffer from high complexity. A simplified ML implementation consists of three parallel search engines, one for each stream. Each engine is comprised of a QR decompositor which decomposes matrixes into orthogonal and right triangular matrixes and an ML detector which intelligently searches for candidates with minimum error distances. In this way, the system is able to dynamically adjust the input into the ML decoder to achieve higher accuracy demodulation (see Figure 3).

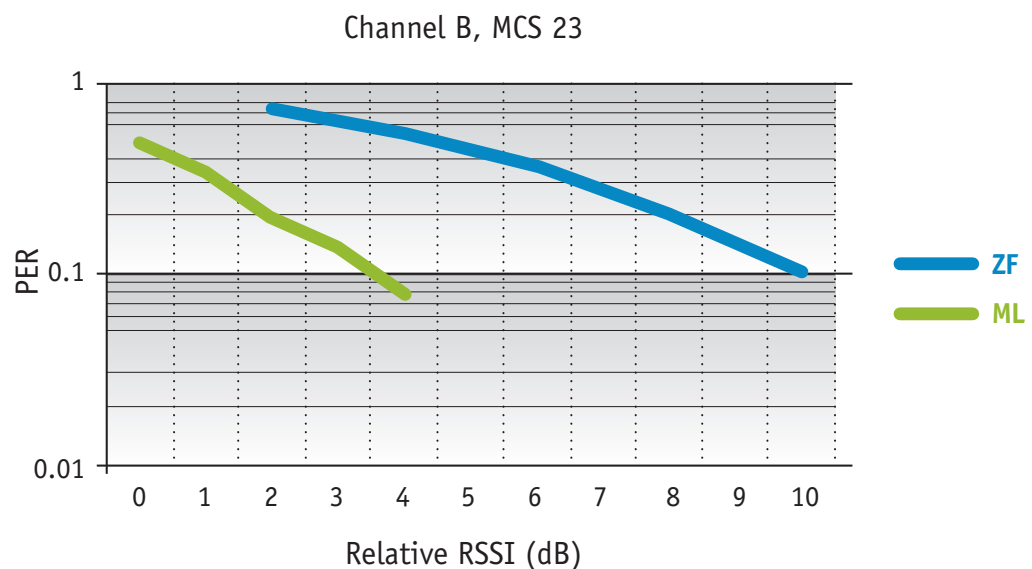


Figure 3: Maximum Likelihood (ML) Equalizer Performance Gain

A Simplified Maximum Likelihood (SML) architecture consists of a QR decompositor and an ML detector which intelligently search for candidates with minimum error distances to achieve higher accuracy demodulation compared to Zero Forcing (ZF) equalization.

SUSTAINING SIGNALS THROUGH THE ENTIRE SIGNAL CHAIN

While supporting Low Density Parity Check (LDPC) codes, Transmit Beamforming (TxBF), and Maximum Likelihood (ML) technologies significantly improve signal quality, signals must be sustained across the entire signal chain to maintain the highest performance and reliability in 802.11n systems. To prevent signal quality degradation and achieve the highest possible signal integrity while controlling system cost, simplifying design, and accelerating product development, developers will need to rely upon a variety of other technologies as well:

- RF front-end integration: integrating both 2.4 and 5 GHz low-noise amplifiers (LNA) and power amplifiers (PA) reduces component count and board size. Estimated LNA noise is 4.6 dB at 2.4 GHz and 5-6 dB at 5 GHz
- PA Linearization: As PAs become nonlinear at large output levels, they cause EVM degradation and spectral growth. PA linearization (PAL) uses a feedback loop built around the PA as well as PA Predistortion (PAPRD) techniques to reduce AM-to-AM distortion, resulting in an ~3.5 dB improvement in output power (see Figure 4)

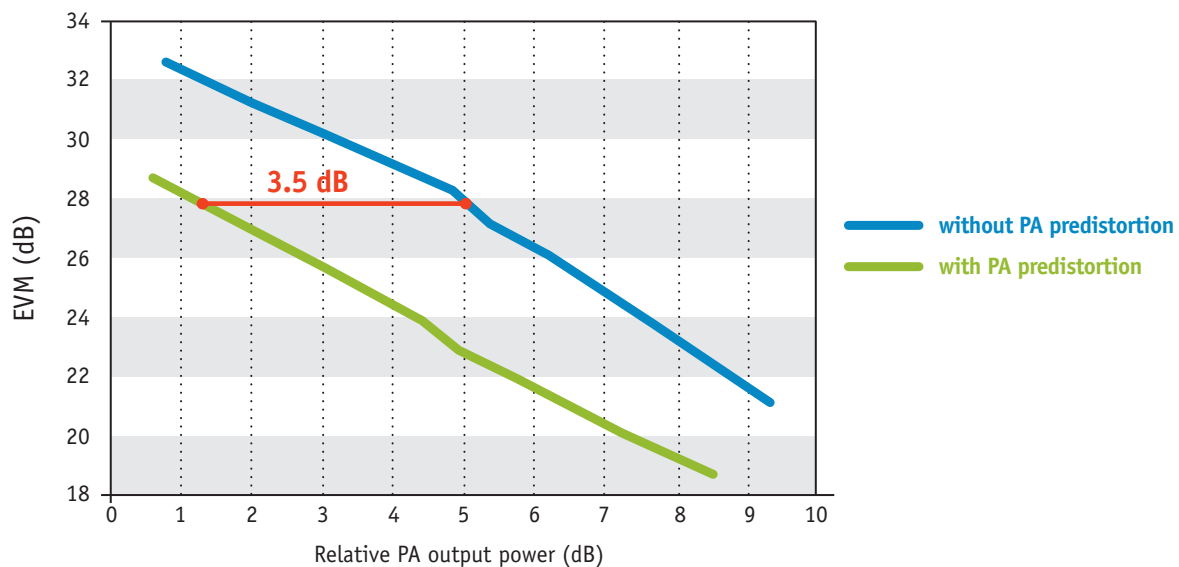


Figure 4: PA Predistortion Improvements

Power amplifiers become nonlinear at large output levels, leading to EVM degradation and spectral growth. Power amplifier predistortion (PAPRD) techniques reduce AM-to-AM distortion, resulting in a ~3.5 dB improvement in output power depending upon the particular implementation.

- Rx Peak Detection: On-chip peak detectors placed after the LNA and before the mixer in all receive chains allows for the detection of front-end circuit saturation arising from out-of-band strong interference. Saturation exceeding a specific threshold will trigger an RF gain drop. This provides for better performance in environments with a high level of interference.
- Space Time Block Coding (STBC): STBC provides added robustness for devices operating in an environment with many reflections and is commonly used in systems where there are multiple transmit chains but only a single receiver chain, such as a handheld device with a single antenna talking to an Access Point (AP). With STBC each AP transmitter, for example, sends a known linear combination of the other transmitters' signals which the single receiver uses to identify the original data stream. Uplink transmissions to the AP also achieve the same signal quality improvement through maximal ratio combining at the AP.

At the AP, each transmitter sends a known linear combination of each other's signals and the single receiver can algebraically determine what the original data stream is, based on prior knowledge of each transmitter's linear combination formula. The asymmetry in the antenna configurations may suggest that STBC only improves range in the downlink direction (AP to client). However, the same improvement in range in the uplink transmission can be attained by utilizing maximal ratio combining at the AP.

- Improved Transmit Power Control: New RF transmit power control methods reduce system complexity by eliminating the need for an off-chip RF-peak detector and its calibration at manufacturing time
- I/Q Calibration: Advanced I/Q calibration techniques achieve a higher error vector magnitude (EVM) by addressing transmit I/Q mismatching. Further improvement is possible when transmit I/Q mismatch is separated from receiver I/Q mismatch using a switchable phase shift in the loopback
- Radio built-in self test (RBIST): Self-test facilitates troubleshooting as well as reduces the time required to complete comprehensive testing during manufacturing processes
- Integrated 5/10 MHz BB channel filter: This filter is used in Enterprise applications and can be implemented by cutting the clock rate with no negative impact on die size
- Revised LO generation using a frequency doubler: Lab measurements show frequency synthesizers using a frequency doubler provide 2 dB phase noise improvement

Signal quality is also directly affected by various characteristics of the 802.11n radio. Reducing internal phase noise, for example, provides a few dB better performance compared to previous generation radios. More backoff in the DAC and improved Transmit baseband/filter linearity, as well as crystal interference mitigation and a selectable PLL clock edge add further gain improvement.

As important as improving performance is, it is critical to make sure that these hard-earned gains are not given up elsewhere in the signal chain. For example, performance gains can be lost if there is no way to get data off-chip efficiently enough to match 802.11n's high data rate. Consider that without packet aggregation, the 450 Mbps PHY data rate results in 56 Mbps throughput. This figure assumes no collisions, and TCP throughput would be even less. Packet aggregation increases throughput, especially with smaller packets, and requires efficient DMA of descriptors and packet data.

To avoid throughput bottlenecks, 802.11n radios require a high-speed internal bus. The newer ARM AXI bus, for example, was designed to enable more efficient DMA transactions compared to the AHB bus. AXI supports split transactions, permits multiple simultaneous memory fetches, and eliminates AHB's requirement that each transaction must complete before the next can be issued. As a result, smaller frames are processed more quickly, software efficiency is improved, and less load is placed on the host, resulting in the unimpeded flow of data to the main application processor.

COST AND PERFORMANCE

The combined impact of implementing all of these improvements in concert is to create a system that provides the highest possible performance. However, to understand the true impact of sustaining the signal and improving reliability requires measurement of actual throughput. For instance, the increase in coding gain realized from LDPC technology improves performance not only by increasing coding efficiency but also by decreasing the BER which, as importantly, reduces the number of retransmissions required. Much the same way one can meet any budget by cutting spending or eliminating an expense, reducing retransmissions results in more efficient utilization of available throughput.

The impact on cost of sustaining the signal can be illustrated by analyzing the cost delta between increases in throughput. Moving to a 3-stream implementation provides three times the bandwidth compared to a single stream. This also requires three times the receivers and transmitters, as well as additional complexity – and all associated costs – to handle synchronization and demodulation. Utilizing leading-edge LDPC, TxBF, and ML technologies, however, effective throughput can be increased by up to 66% or range extended up to 50% further. With the availability of highly-integrated SoC architectures, this added throughput and range comes at a nominal cost relative to the total system cost.

802.11n is a complex technology presenting many design challenges that affect signal quality. By sustaining signals throughout the signal chain to maintain signal integrity and reliability, it becomes possible to design efficient systems which achieve higher throughput and greater range. In this way, engineers can create the highest-performing, most reliable, and cost-effective 802.11n systems on the market.