Making OFDM Work for High-Performance Wireless Network Applications

James A. Crawford, CTO & Co-Founder
Magis Networks, Inc.

1. Introduction

Automobiles have been with us now for many decades and in a general sense they are all very similar, all requiring 4 wheels, an engine and fuel system, a chassis, etc. However, no one would mistake a Ferrari for a Jeep because their intended use, design and costs differ greatly. In the same way, wireless networks for the home and office have been with us for at least the past 10 years, but delivered performance can differ greatly depending upon the objectives in mind.

The most dominant wireless networking standard of the day is IEEE802.11b which was originally designed approximately 10 years ago for data-only based communications. In order to make the networking protocol as simple and scalable as possible, a carrier-sense-multiple-access (CSMA) medium-access control (MAC) protocol was adopted. With the advent of a wide spectrum of new consumer electronic (CE) applications like high-definition television (HDTV), MP3 players, digital cameras and camcorders, personal video recorders (PVRs), cable modems, 1394 applications, satellite television, etc., the limited capabilities offered by 802.11b networks have been dramatically surpassed and newer wireless technology is greatly needed.

Enter Air5™, an OFDM-based wireless networking technology that offers the information throughput rate, communication reliability and quality of service (QoS) that is needed by this host of new CE devices and applications. Comparing Air5 to existing IEEE802.11b networks is not that different than the earlier automobile comparison, and in the balance of this article, some of these important differences will be made more clear.

Perspective has been very key in the development of Air5. Magis' lead investors mandated in mid-1999 when the start-up company was founded that the wireless technology developed be capable of "delivering multiple streams of high-quality video simultaneously" within the home and office along with other less demanding services like voice and data. At the same time, the technology had to be affordable if it had any chance of succeeding in the CE marketplace, and standards would also be an important issue. Although the IEEE802.11a wireless standard had been released in the late 1990's, careful technical assessment of the standard showed that it alone would be unable to deliver the level of performance needed. Rather than completely abandon IEEE802.11a for an entirely new system design however, it was found that improvements could be made to the OFDM baseline system design that would permit the modified system to achieve the needed objectives and Air5 was consequently born. Some of the factors that make Air5 the unprecedented superior wireless vehicle for delivery of video, voice and data are discussed in the sections that follow.

It should be understood from the outset that high-QoS video delivery is the most demanding of all possible services, and that although many references to the video case will be mentioned in this memorandum, this in no way should convey that Air5 is meant for video alone. Rather, the same high QoS and throughput necessary for video benefits the wireless delivery of voice and data as well.

This paper primarily addresses the physical layer (PHY) and medium access control (MAC) layer innovations and considerations that were developed in order to make the OFDM-based system work and work extremely well at 5 GHz.

2. Misconceptions Regarding the 5 GHz Wireless Channel

Many different opinions have been offered about propagation over the 5 GHz channel in recent months. Some quantitative insight into this important question is offered in the following sections.

2.1 Propagation Differences at 5GHz Compared to 2.4GHz

It is common practice to estimate propagation loss over a free-space channel using the Friis[1] formula which gives the receive signal-to-noise ratio (SNR) as

$$SNR_{dB} = 10Log\left[P_T G_T G_R \left(\frac{\lambda}{4\pi R}\right)^2\right] - 10Log\left(kT\ BW\ NF\right) \quad (1)$$

where

P_T	Transmit power in Watts
G_T	Numerical transmit antenna gain[2]
G_R	Numerical receive antenna gain
λ	Wavelength, meters
R	Range, meters
k	Boltzmann constant
T	Absolute temperature, taken to be 290K
BW	Bandwidth of modulation, Hz
NF	Noise factor of receiver[3]

For use with indoor communications where flat losses due to walls and other materials are present along with multipath-related issues, this equation is normally modified to

$$SNR_{dB} = 10Log\left[P_T G_T G_R \left(\frac{\lambda}{4\pi}\right)^2\right] - 10nLog(R) - 10Log(kT\ BW\ NF) - L_{dB} \quad (2)$$

where the "range loss exponent" is given by n, and L_{dB} is the bulk loss (in dB) due to absorption by walls, etc. For free-space propagation, n≡2 and L_{dB} is 0 dB. The primary quantities of interest for the indoor channel are of course the loss exponent n and the bulk loss L_{dB}.

A second straight-forward model that has been considered from time to time within Magis is that by Medbo [2]. This model assumes an additional flat dB-per-meter loss represented by α in (3) but is otherwise the free-space model of Friis.

$$SNR_{dB} = 10Log\left[P_T G_T G_R \left(\frac{\lambda}{4\pi R}\right)^2\right] - 10Log(kT\ BW\ NF) - \alpha R \quad (3)$$

In (3), R is the range in meters and α has the units of dB / meter having a value of approximately 0.44 dB / m.

All of the channel propagation loss models just mentioned exhibit a square-law decrease in SNR as the RF frequency is increased. If all other factors are left unchanged in the foregoing equations, the SNR for a signal at 5.25 GHz rather than 2.4 GHz will be reduced by $20Log(5.25/2.4) = 6.8\ dB$. If we were dealing with a strictly line-of-sight (LOS) channel without multipath using ***omni-directional antennas***, the discussion would be over. However, for a LOS-link at 5 GHz, this 6.8 dB difference could be easily made up by the increased antenna gain available at 5.25 GHz as compared to 2.4 GHz with the same size antenna. The presence of significant multipath with the indoor wireless channel changes the picture dramatically though, especially as the data rate is increased beyond several Mbps, and these simplistic models are inadequate to reveal the entire story.

2.1.1 Propagation Losses Through Common Building Materials at 2.4 GHz and 5 GHz[4]

Additional propagation losses beyond that experienced in free-space are due primarily to (i) reflections caused by spatial impedance discontinuities encountered by the signal wavefront as it propagates through space and (ii) ohmic losses that occur from propagation through materials other

[1] D.C. Hogg, "Fun with the Friis Free-Space Transmission Formula", IEEE Antennas and Propagation Magazine, Vol. 35, No. 4, August 1993
G.W. Collins, "Wireless Wave Propagation", Microwave Journal, July 1998

[2] Numerical Gain= $10^{0.1 Gain_{dB}}$

[3] Noise Factor = $10^{0.1 NoiseFigure_{dB}}$

[4] Magis Networks Report, [3]

than free-space. Since the impedance of free space is $Z_o = \sqrt{\mu_o / \varepsilon_o} = 377\,\Omega$, any impedance change due to walls, floors, etc. encountered by the propagating signal will result in reflections that lead to frequency-selective fading due to the multipath.

For measuring the scattering parameters of a dielectric material in free space, the relative permittivity of the material can be calculated as described here by assuming a planar incident wavefront and an infinite plane-parallel plate dielectric slab. Imposing boundary conditions at the interface of the dielectric material, that is, that the tangential components of the electric and magnetic fields must be continuous, leads to a system of equations relating the transmission and reflection coefficients of the system, the electric fields, and the dielectric properties of the material. These quantities are described graphically in Figure 1 and the equations that follow

Figure 1 Electric Field Components for a Plane Electromagnetic Wave Incident on an Infinite Plane Dielectric Slab in Free-Space

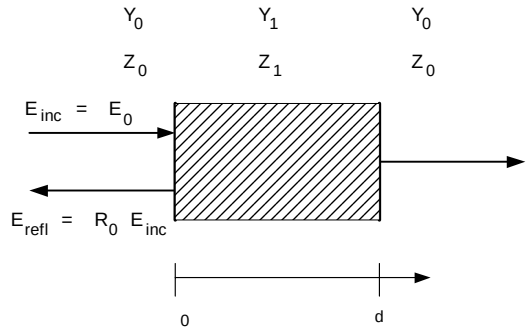

$$
\begin{aligned}
&E_0(1+R_0) = E_1(1+R_1)\\
&\frac{E_0}{Z_0}(1-R_0) = \frac{E_1}{Z_1}(1-R_1)\\
&E_1\left(e^{\gamma_1 d} + R_1 e^{-\gamma_1 d}\right) = T_0 E_0\\
&\frac{E_1}{Z_1}\left(e^{\gamma_1 d} - R_1 e^{-\gamma_1 d}\right) = T_0 \frac{E_0}{Z_0}
\end{aligned}
\qquad (4)
$$

where $\gamma_i = j(2\pi/\lambda_0)\sqrt{\tilde{\varepsilon}_{ri}}$, $Z_i = \frac{\gamma_0}{\gamma_1} Z_0$, $\tilde{\varepsilon}_{ri}$ is the complex relative permittivity of medium "i" and Z_i is the impedance of medium "i". The solutions in terms of R_0 and T_0 for this system of equations are given in [1] as

$$
\begin{aligned}
R_0 &= \frac{\left(\gamma_0^2 - \gamma_1^2\right)e^{-\gamma_1 d} - \left(\gamma_0^2 - \gamma_1^2\right)e^{\gamma_1 d}}{\left(\gamma_0 + \gamma_1\right)^2 e^{-\gamma_1 d} - \left(\gamma_0 - \gamma_1\right)^2 e^{\gamma_1 d}}\\
T_0 &= \frac{4\gamma_0\gamma_1}{\left(\gamma_0 + \gamma_1\right)^2 e^{-\gamma_1 d} - \left(\gamma_0 - \gamma_1\right)^2 e^{\gamma_1 d}}
\end{aligned}
\qquad (5)
$$

The power transmission and reflection coefficients are then given by

$$
\begin{aligned}
T &= \left|T_0^2\right|\\
R &= \left|R_0\right|^2
\end{aligned}
\qquad (6)
$$

and $T + R + A = 1$, where A is the power coefficient of absorption.

Representative transmission and reflection results from [3] are provided here in Table 1. As shown there, most of the losses are very similar between 2.3 GHz and 5.25 GHz except for red brick and cinder block where the 5.25 GHz signal losses are higher.

Table 1 Transmission and Reflection Coefficients at 2.3GHz and 5.25GHz [3]

Material	T (dB)			R (dB)		
	2.3 GHz	5.25 GHz	Δ	2.3 GHz	5.25 GHz	Δ
Plexiglass (7.1mm)	-0.3560	-0.9267	0.5707	-12.23	-5.65	-6.5753
Plexiglass (2.5mm)	-0.0046	-0.2041	0.1994	-21.69	-13.25	-8.4470
Blinds (closed)	-0.0016	0.0020	-0.0035	-30.97	-20.39	-10.578
Blinds (open)	0.0137	0.0315	-0.0178	-44.23	-46.95	2.7210
Red brick (dry)	-4.4349	-14.621	10.186	-12.53	-8.98	-3.5459
Red brick (wet)	-4.5119	-14.599	10.087	-12.52	-9.41	-3.1185
Carpet (back)	-0.0361	-0.0318	-0.0044	-25.19	-15.8	-9.4080
Carpet (weave)	-0.0271	-0.0056	-0.0214	-26.94	-18.7	-8.2710
Ceiling tile	-0.0872	-0.1795	0.0923	-21.07	-18.7	-2.3470
Fabric	0.0216	0.0133	0.0083	-41.70	-30.1	-11.570
Fiber-glass	-0.0241	-0.0340	0.0099	-39.40	-28.8	-10.581
Glass	-0.4998	-1.6906	1.1908	-11.29	-4.9	-6.3446
Drywall (12.8mm)	-0.4937	-0.5149	0.0211	-12.11	-11.5	-0.6390
Drywall (9mm)	-0.5095	-0.8470	0.3376	-12.03	-8.87	-3.1596
Light cover (front)	-0.0040	-0.0533	0.0494	-28.47	-20.0	-8.4490

Material	T (dB)			R (dB)		
	2.3 GHz	5.25 GHz	Δ	2.3 GHz	5.25 GHz	Δ
Light cover (back)	-0.0070	-0.0532	0.0462	-28.07	-18.8	-9.2390
Linoleum (back)	-0.0186	-0.1164	0.0977	-26.05	-17.3	-8.7610
Linoleum (front)	-0.0198	-0.1278	0.1081	-23.69	-16.0	-7.6690
Fir lumber	-2.7889	-6.1253	3.3364	-17.45	-14.8	-2.6890
Particle Board	-1.6511	-1.9508	0.2997	-8.59	-14.1	5.5359
Plywood	-1.9138	-1.8337	-0.0801	-9.05	-30.5	21.422
Stucco (back)	-14.582	-13.906	-0.6760	0.62	0.04	0.5785
Stucco (front)	-14.863	-13.235	-1.6280	-2.38	-9.24	6.8587
Tiles	-2.2163	-1.4217	-0.7946	-6.24	-14.9	8.6093
Tar paper	-0.0956	-0.1341	0.0385	-28.88	-17.8	-11.067
Cinder block (dry)	-6.7141	-10.326	3.6119	-7.67	-6.13	-1.5324
Cinder block (wet)	-7.3527	-12.384	5.0313	-5.05	-7.55	2.5080
Diamond mesh	-20.985	-13.165	-7.8200	-0.53	0.89	-1.4216
Wire lath (paper)	-1.2072	-0.7044	-0.5028	-6.38	-10.9	4.6015
Wire lath	-1.2136	-0.3404	-0.8732	-8.01	-21.8	13.764

It is worthwhile to point out that signals reflected from many of the materials including glass and plexiglass in Table 1 are less attenuated at 5 GHz than at 2.4 GHz. This fact translates into richer multipath characteristics at 5 GHz than at 2.4 GHz which can be advantageously exploited to improve non-LOS links within indoor environments if desired.

Key Point: *Propagation losses through most building construction materials are very similar at 2.3 GHz and 5.25 GHz. Reflections are however more prevalent at 5 GHz leading to richer multipath characteristics that if advantageously exploited, can lead to very substantial improvements in link throughput and reliability.*

2.2 Time-Varying Multipath

The short wavelengths involved at both 2.4 GHz and 5.25 GHz lead to multipath characteristics that can vary fairly quickly with time. At the same time, nearly time-static multipath fading characteristics have also been observed that span seconds or even minutes. This latter phenomenon makes the use of large data buffers to introduce time diversity with standard IEEE802.11a implementations questionable at best because even very long buffers (e.g., 90 seconds) can be over-run. The issue being addressed within this section pertains more to the question of how quickly the channel multipath characteristics can change rather than how long they may persist.

In order to facilitate this discussion, it is helpful to consider a simple 3-ray multipath model that can be represented mathematically with a channel transfer function given by

$$H(f) = 1 + A_1 e^{-j2\pi f(\tau_1+\tau_m)\tau_1} + A_2 e^{-j2\pi f(\tau_2+\tau_m)} \quad (7)$$

where A_1 and A_2 are the relative strengths of the two multipath rays, τ_1 and τ_2 are the relative path length time differences compared to the direct-ray, and τ_m is a hypothesized additional time-dependent path delay due to user movement. As this model stands, it can exhibit channel loss or channel gain above the normal free-space channel path loss. Selection of different parameter values is fairly non-critical in showing the time-varying behavior of the multipath characteristics that result.

In free-space, signals travel approximately 11.8 inches per nanosecond. If the user movement is assumed to be V_{user} in feet-per-second orthogonal to the receive signal wavefront, then $\tau_m(t) = 12V_{user}t/11.8$ nsec. The sensitivity of the observed frequency-selective fading shown in Figure 2 to path delay changes in (7) is incredible in that even 0.1 nsec represents 1.18 inches which is almost a half-wavelength. In the case of Figure 2, the example parameter values chosen were A_1= 1, A_2= 0.40, A_3= 0.23, τ_1= 21.85 nsec, τ_2= 32.34 nsec and τ_3= 53 nsec and the user was assumed to be moving at 5 feet per second which is equivalent to about only 3 miles per hour.

Figure 2 Example Channel Multipath Characteristics Over Time

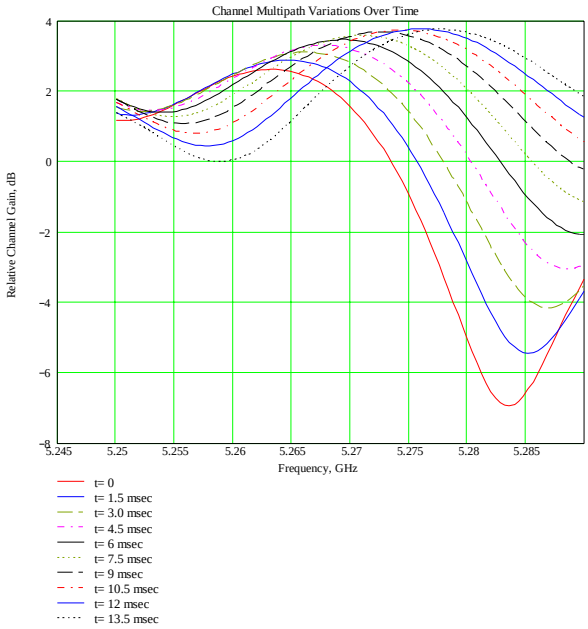

As shown in this simple example, a small amount of motion introduced into the wireless channel can result in a wide 10 dB frequency-selective impairment that could break the wireless link along the edge of coverage.

In order to explore the channel coherence issues involved with communication over the 5 GHz channel, the Air5 system was purposely modified to use longer MAC frames over several different channel conditions and histograms of the data packet error rate (PER) were computed. The channel estimates that were used were based upon single-shot estimates using the T1-T2 long-symbol portion of the IEEE802.11a standard preamble. The channel that was considered was the large conference room area at Magis which is a large open-area having large metallic window blinds along two walls, an approximately square perimeter having about 1500 square feet of floor space, 8 foot ceiling in an industrial multi-story steel-reinforced concrete building. The wireless link was set up across a distance of about 25 feet in an area to the side of the audience seating area where refreshments were served during a company meeting. The "channel 1" conditions occurred while approximately 40 people were milling around getting refreshments and directly disturbing the line-of-sight link. The "channel 2" conditions are actually a long-term average over the entire length of the meeting that was held in the conference room thereby including a substantial amount of time during which most of the staff was seated. The measurement data was post-processed to compute the PER as a function of the cell location within each MAC frame and the results are shown in Figure 3. As fully expected, the "channel 1" conditions with many people moving within the conference room exhibited considerably worse PER than the more benign conditions that involved far less people activity on the average.

The point of this channel coherence discussion is primarily that the operating point of the entire system must be managed carefully in order to deliver good QoS performance. Good QoS performance without using large data buffers and their accompanying time delay limits the number of packet re-transmission attempts that can be used. If the PER is too high and the number of re-transmissions for a given packet are exhausted, that packet would be dropped and artifacts introduced if operating in video mode. The results here are specific to the implementation details of the Air5 system, but the same kind of factors must be considered in any other system design also.

Figure 3 Packet Error Rate Versus Packet Location Within MAC Frame for 64-QAM Rate= ¾

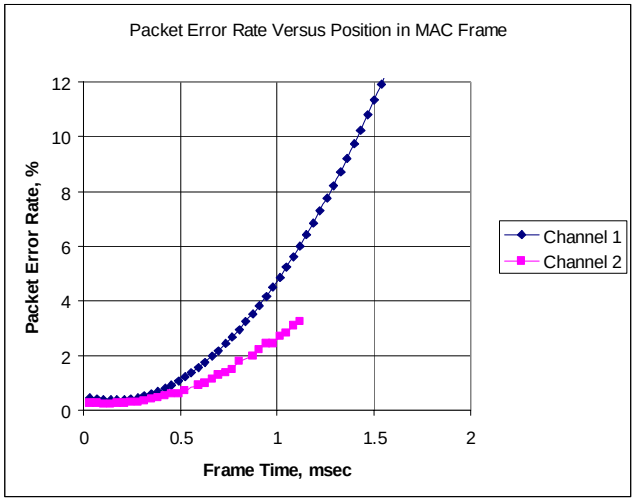

Several summary points are worth noting with regard to the time-varying nature of the frequency-selective fading / multipath that is present on most 5 GHz indoor channels:

Key Points:

- *With a free-space wavelength of only about 2 inches, very slight positional displacements can change the frequency-selective fading characteristics dramatically;*
- *Even very modest movement of either the wireless terminals or other objects within the propagation volume can lead to significant changes in the frequency-selective fading characteristics;*
- *Without some form of channel adaptation, use of long MAC frames will be very difficult due to channel coherence losses, particularly at 36 Mbps and higher. The PER must be closely administrated if high QoS performance is desired.*

2.2.1 Pedestrian-Induced Fading at 5 GHz

A recent paper considered the effects of pedestrian traffic on indoor 5 GHz channels [5] ***based upon simulation***. Fading profiles were reported for a point-to-point link in a 150-m^2 open-plan environment with moderate pedestrian traffic conditions. The fading depths reported ranged from 31 dB to 36 dB and the results were Rayleigh-distributed despite the presence of a direct-ray for the majority of the simulated scenarios. The floor plan used for the simulation study is shown here in Figure 4. The pedestrian-induced fading episodes that this simulation study predicted are shown in Table 2. Many factors are not included in this simulation model including the modulation bandwidth and the role of frequency-selective fading versus a complete flat fade across the entire modulation bandwidth. In this respect, experimental results using the IEEE802.11a waveform would probably be more useful. Even so, if the below-threshold percentages that are reported in Table 2 are representative of the channel behavior involved, this level of performance may be quite adequate for data-only applications whereas it is not acceptable for the high QoS demands posed by video.

Figure 4 Floor Plan for Simulation Study Regarding Pedestrian Traffic at 5 GHz

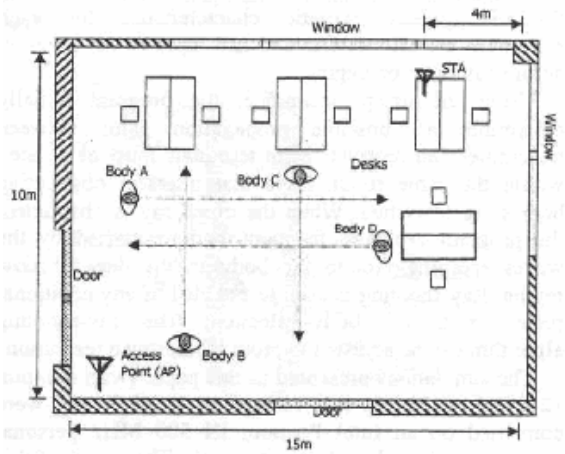

Table 2 Pedestrian-Induced Fading Episodes at 5.7 GHz

Threshold Level Relative to Mean (dB)	% of Time Below Threshold
0	66.6
-5	21.1
-10	5.8
-15	1.1
-20	0.1

Key Point: *If the below-threshold probabilities are viewed as video-outage probabilities, pedestrian traffic must be a major consideration in the system design in the context of high-quality video delivery. Short of using very large video buffers, no one wants to receive only 94.2% of their 2 hour HDTV movie error free.*

2.2.2 Generalized Space-Time Processing to Combat Channel Multipath

Channel delay-spread associated with channel multipath often limits the range performance of IEEE802.11b systems rather than inadequate signal strength, particularly at the higher data throughput rates. Since multipath is really interference of one's own signal with oneself, severe multipath can actually be more problematic for smaller distances than for larger distances.

Space-time coding (STC) and its many variants is one of the most active research areas in wireless communications. Generally, the primary motivation for STC is to increase the overall system capacity in terms of bits/Hz. Unfortunately, these systems are still quite expensive and in the most degenerate channel cases, very limited capacity improvement may be achieved.

STC systems take advantage of the transmit and receive signal wavefronts in the time, spatial and coding dimensions to ideally deliver greater

throughput. A far less aggressive perspective has been employed for many years through traditional diversity methods which were developed in part to improve channel reliability in the wake of channel multipath. Traditional diversity systems employ a wide range of techniques including spatial, frequency, time and polarization methods in order to combat channel multipath as well as other channel impairments.

Most if not all of the primary channel loss models that are used throughout the wireless industry are related in some way to a power-law range dependency as discussed earlier in Section 2.1. All of these formula rely on local CW signal strength measurements that are made over a specified spatial region and then averaged. This time and position averaging eliminates almost all of the spatial and frequency domain structure present. Take for instance the discussion of the measurement techniques used in a recent paper[5] on this topic:

"Halfwave dipole (2 dBi) antennas and directional patch antennas (7 dBi, 90° x 70° beamwidth) were used in the measurements. The transmitted signal was a continuous wave at 5.2 GHz and of about 30 dBm power. At the receiver, a spectrum analyzer (HP8595E) with a low noise amplifier was used. The sensitivity[6] of the receiver chain was about -130 dBm. For each measurement point a spatial average of received power was obtained by taking the median of 401 samples on a horizontal circle of about 0.5 m diameter."

As clearly stated, spatial signal strength variations of the signal wavefront are purposely averaged out. Therefore, if one specific spatial location exhibits very good signal strength, even though the system could benefit greatly if it somehow chose to receive the signal at that location, the power-law assessment averages all of this detail away. Furthermore, our own experience at Magis has shown that the spatial profile of the signal strength can change dramatically unless "all" of the measurements are made very nearly at the same instant in time.

Many channel sounding assessments were done in the Magis building two years ago to more fully understand the propagation problem more thoroughly. A typical result from one of these assessments is the time-frequency domain chart provided here in Figure 5. During this effort, a wide-band CDMA signal source was used having a modulation bandwidth of about 20 MHz. Signal interception software was developed in a laptop computer environment for ease of mobility and a pair of simple half-wave patch antenna were used to capture the signal wavefront at two points in space simultaneously. As shown in Figure 5, the receive signal structure is clearly time-varying with many rich features. In each plot, two traces are overlaid corresponding to the signal received by the two different patch antennas separated a specified distance apart.

Figure 5 Example Receive Signal Assessment where Signal Amplitude is Plotted Against Time and Frequency

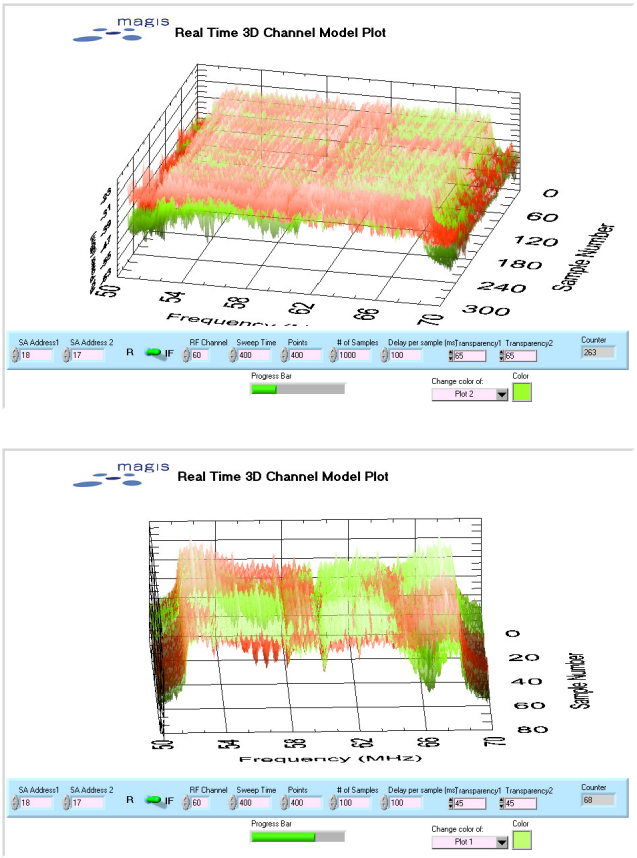

Although colorful, this previous perspective of the time-varying frequency-selective fading channel is not sufficiently quantitative to be that useful. A second-generation channel instrumentation setup was put together with the collection patch antenna hosted on a pair of plastic micrometers as shown in Figure 6. Care was taken to make sure that

[5] Medbo, J., Jan-Erik Berg, "Simple and Accurate Path Loss Modeling at 5 GHz in Indoor Environments with Corridors", VTC2000

[6] Although not stated, if we assume that the measurement system noise figure is 8 dB, this translates into a resolution bandwidth of about 4 kHz.

the channels involved were relatively time-stationary. As evidenced in Figure 7 through Figure 11, a very slight horizontal re-positioning of the collection antenna could changed the frequency-selective fading characteristics dramatically.

Historically speaking, there seems to be a severe absence of publicly available literature that reveals the volatility of the indoor 5 GHz channel. There is nothing said in the IEEE802.11a standard about this topic and all of the performance requirements are posed in the context of a classical AWGN channel. No criticism is meant by this statement because the Magis team has also struggled to quantitatively describe the complexities of this channel for our own purposes including computer simulations; the task is far from easy or straight forward. In the end, we have relied heavily on divide-and-conquer analysis methods to deal with the channel complexities. The Magis staff have often quipped that "the perfect Ph.D. dissertation would be one on propagation at 5 GHz because (i) no one could ever disprove your work or (ii) repeat it for that matter!"

Considerably more material is provided on this all-important topic in Section 4.

Key Point: *Power-law based propagation models are useful for a very high-level perspective of the indoor propagation problem but they average out all of the multipath features that can in principle be exploited to dramatically improve system performance.*

2.2.3 Outage Probability

The figures and discussion presented thus far were in part intended to illustrate how easily multipath (with a single-antenna system) can lead to severe frequency-selective fading. The time duration of a particular fading characteristic can be several msec or even many seconds. So unlike the mobile wireless channel where user-mobility virtually ensures that severe fades will not last overly long, the potential for long-lasting poor channel conditions on the indoor 5 GHz channel is high, particularly in the context of a 2 hour length HDTV movie.

It is well understood that OFDM inherently offers a great deal of immunity to frequency-selective fading in principle. However, any realistic radio can only deliver a limited amount of SNR, being inherently limited by such factors as (i) the phase noise of the local oscillator, (ii) linearity of the transmitter power amplifier (PA) and (iii) linearity of the receiver. In the context of 64-QAM rate ¾ coding which is used for the highest IEEE802.11a mode (54 Mbps), a minimum SNR for an AWGN channel of roughly 25 dB is required whereas the receiver can deliver perhaps 30-35 dB under the best of conditions[7]. Consequentially, it is impossible to handle appreciable frequency-selective fades of more than 5-10 dB regardless of the signal power received.

Figure 6 Patch Antenna Configured on Two Plastic Micrometers for Precision Spatial Positioning

Figure 7 20 MHz Frequency Sweep Centered at 5.23 GHz, X= 0.20, Y= 0.0 (in.)

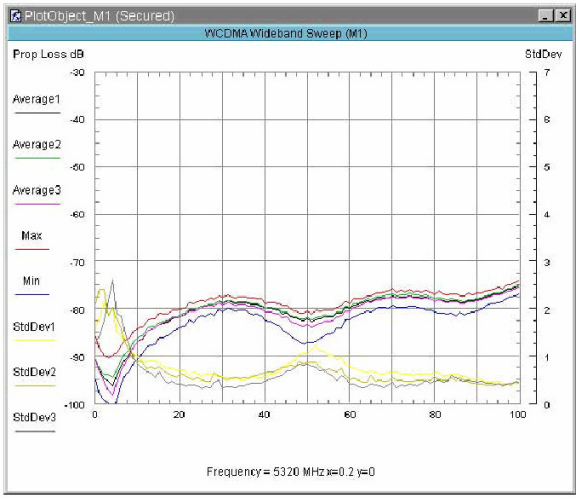

[7] For example, total-link (i.e., transmit plus receive) local oscillator phase noise of 1 degree rms limits the SNR to about 38 dB alone.

Figure 8 20 MHz Frequency Sweep Centered at 5.23 GHz, X= 0.40, Y= 0.0 (in.)

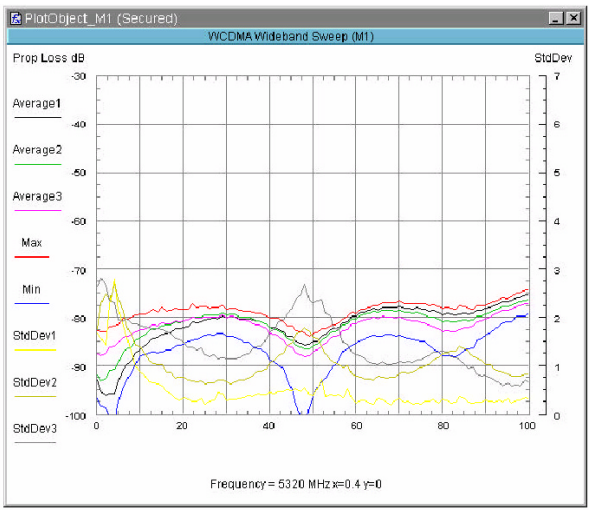

Figure 9 20 MHz Frequency Sweep Centered at 5.23 GHz, X= 0.60, Y= 0.0 (in.)

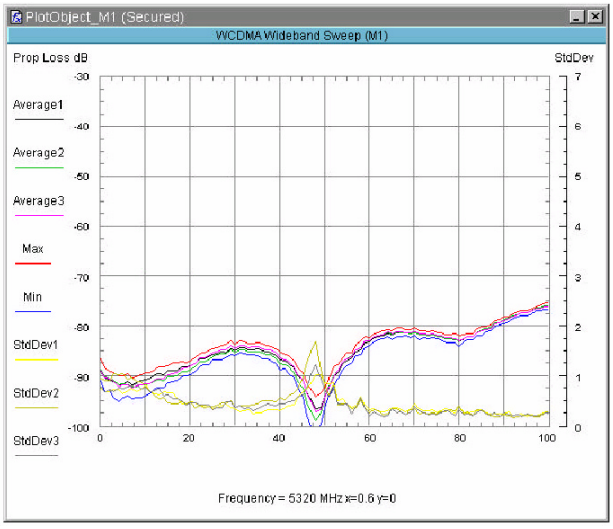

Whether the spot SNR is limited by AWGN combined with frequency-selective fading or limited by phase noise and radio nonlinearity contributions, this previous discussion should show that it is fairly easy to experience link outages with a single-antenna system during which QoS constraints are not achieved. If the probability of such an outage each second is p_{out} for a single "look" at the receive signal wavefront, multiple independent looks at the receive wavefront can in principle reduce the outage probability significantly. For the purposes of illustration only, if 5 such independent looks at the receive signal wavefront were available with the same individual outage probability p_{out}, the probability that the system still experiences an outage would be roughly p_{out}^5 which is a dramatic improvement over a single-look system. This is obviously a description of what simple diversity systems are intended to accomplish, but the need to deliver low-outage high-reliability channels in the case of high quality video cannot be understated.

Figure 10 20 MHz Frequency Sweep Centered at 5.23 GHz, X= 1.0, Y= 0.0 (in.)

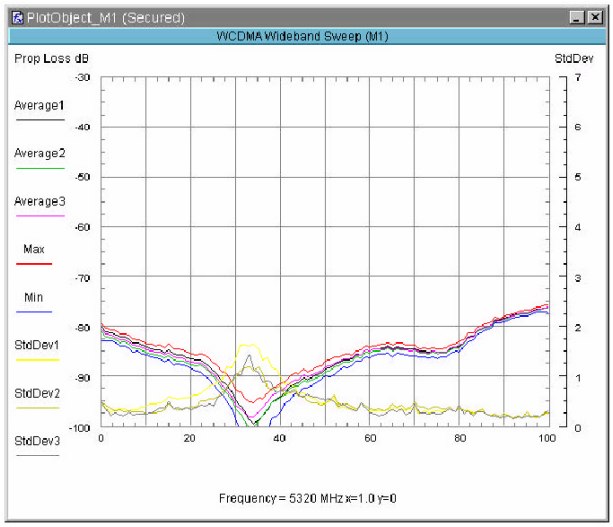

Figure 11 20 MHz Frequency Sweep Centered at 5.23 GHz, X= 1.8, Y= 0.0 (in.)

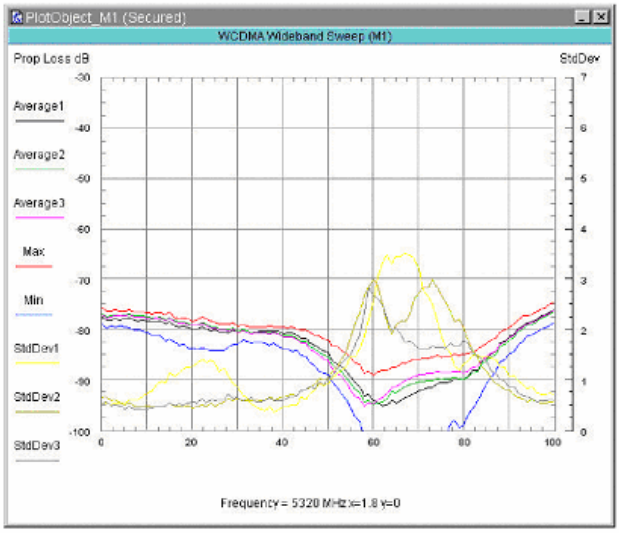

Key Point: *In a single-antenna system, the frequency-selective fading characteristics can change dramatically if the receive antenna is moved as little as one-quarter of an inch. The fade width and depth may be very severe and last many seconds.*

3. Data Communication Requirements in Contrast to MPEG2 Video / Multi-Media Requirements

One of the central themes regarding data versus multi-media communications is unquestionably the quality-of-service (QoS) piece which ties together the simultaneous characteristics that a system must deliver in the way of throughput rate, time latency and jitter, and bit error rate. Generally speaking, a system must be designed from the beginning to deliver the highest QoS performance that will be asked from it rather than attempt to add this on as an after-thought. While the latter is possible to a degree, many inefficiencies will remain and the best-QoS throughput rates will be generally fairly low.

3.1 Data Packet Size and Its Role in QoS

The data packet size used in a system that includes CRC-based ARQ should be chosen carefully based upon the lowest sensitivity threshold at which the system is expected to operate. The average number of transmission attempts required to ultimately deliver the packet to the destination without error and other key performance measures can be analyzed by using the simple state diagram shown below. Assume that the probability of an individual data packet error prior to the inclusion of any ARQ is given by p_{er}. The channel probability of an individual bit error is later denoted by p. The analysis that follows first develops the relationships between bit error rate p, raw packet error rate p_{er}, and optimizing the length of the data packet chosen to maximize the average data throughput when ARQ activity is included.

In the analysis that follows, no limit on the number of ARQ re-transmissions is imposed for simplicity. Furthermore, the ability of the CRC error check mechanism is assumed to be perfect. The transfer function that describes the packet delivery from the source in Figure 12 to its ultimate error-free delivery to the destination can be described by

$$H(z) = (1-p_{er})z + (1-p_{er})p_{er}z^2 + (1-p_{er})p_{er}^2 z^3 \ldots$$
$$= (1-p_{er})\sum_{n=1}^{\infty} z^n p_{er}^{n-1} \qquad (8)$$
$$= \frac{(1-p_{er})z}{1-zp_{er}}$$

Figure 12 State Diagram for PER Computation with ARQ

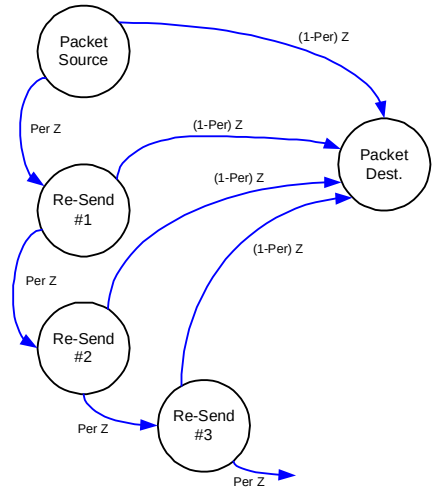

From this representation, it follows that the average number of transmission attempts required to send a given packet successfully to the destination is given by

$$N_{av} = \left(\frac{dH(z)}{dz}\right)_{z=1} = \left(\frac{d}{dz}\frac{(1-p_{er})z}{1-zp_{er}}\right)_{z=1} \qquad (9)$$
$$= \frac{1}{1-p_{er}}$$

The main objective of this analysis is to maximize the net data throughput of the system at a given raw bit error rate (i.e., at the sensitivity point of the system). Let T_b be the time required to transmit a single bit, either data or CRC information. The average time to successfully deliver L_{data} bits to the destination is then

$$T_{av} = \frac{(L_{data} + L_{crc})T_b}{1-p_{er}} \qquad (10)$$

where L_{crc} is the number of CRC bits appended to each data block for error checking. This result can be easily translated to an average bit rate of

$$bps = \frac{(1-p_{er})}{L_{data} + L_{crc}} T_b^{-1} L_{data} \qquad (11)$$

Since successful delivery of the L_{data} bits requires that all of the data and CRC bits be correct, the packet error rate is given by

$$p_{er} = 1-(1-p)^{L_{crc}+L_{data}} \quad (12)$$

where p is the probability of a single bit error. Substituting (12) into (11) results in

$$bps = \frac{(1-p)^{L_{crc}+L_{data}}}{L_{crc}+L_{data}} T_b^{-1} L_{data} \quad (13)$$

The data throughput can be maximized with respect to L_{data} (while assuming that the CRC length L_{crc} is constant) by differentiating (13) with respect to L_{data} and setting the result to zero. Following this process through, the optimal size of the data portion of each packet is the solution given by

$$L_{dataopt} = -\frac{L_{crc}}{2} + \frac{1}{2}\sqrt{L_{crc}^{\ 2} - \frac{4L_{crc}}{\log_e(1-p)}} \quad (14)$$

The resultant $L_{dataopt}$ versus bit error rate p using (14) is shown in Figure 13 with the accompanying resultant packet error rate p_{er} shown in Figure 14.

Figure 13 Optimized Data Packet Length Versus BER per (14) with L_{crc}= 32

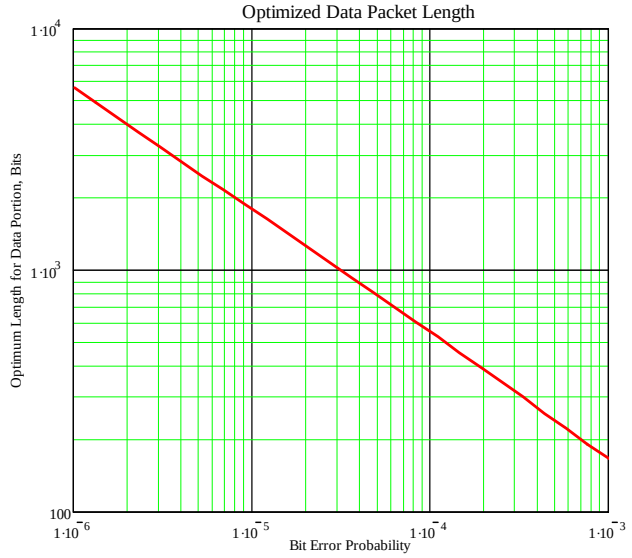

In order to deliver acceptable QoS performance, the number of ARQ attempts (i.e., re-transmission attempts for a specific data packet) must be constrained while also delivering acceptably low BER performance. In this regard, it is convenient to look at the second-moment for the time of arrival of each packet. As made use of earlier in (6)

Figure 14 Packet Error Rate Versus BER Using Optimized Data Packet Length per (14) with L_{crc}= 32

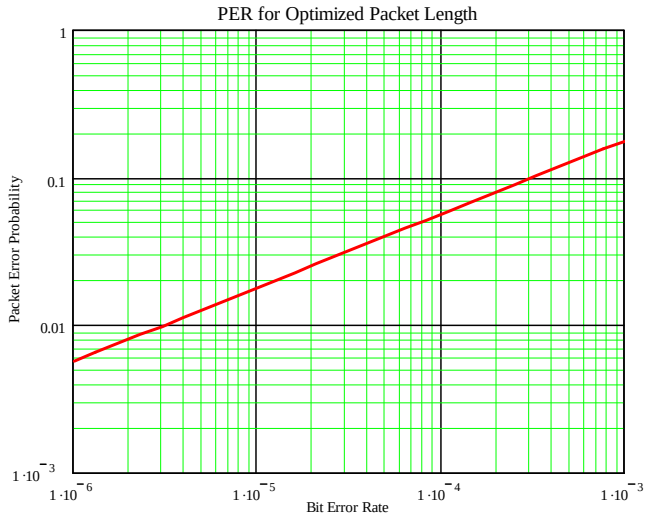

Figure 15 Correct Data Bits Received Versus Total Bits Sent (Including CRC) Versus BER, L_{crc}= 32 Bits

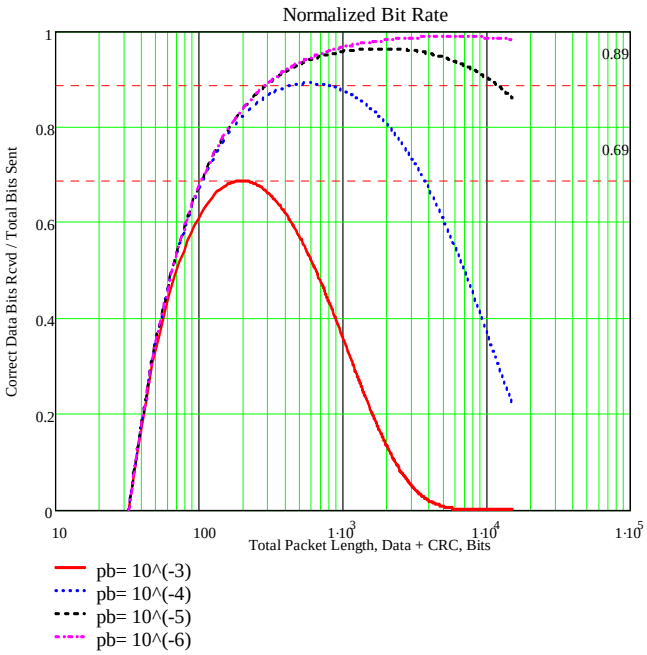

$$N_{av} = \left[\frac{dH(z)}{dz}\right]_{z=1} \quad (15)$$

It can be similarly shown that [4]

$$E\{n^2\} = (1-p_{er})\sum_{n=1}^{\infty} n^2 p_{er}^{n-1} = \left[\frac{d^2H(z)}{dz^2} + z^{-1}\frac{dH(z)}{dz}\right]_{z=1} \quad (16)$$

Using (8) again for H(z) and making use of the rule given by (16), the expected number of transmission attempts rms is given by

$$N_{rms} = \sqrt{\frac{3p_{er} - p_{er}^2}{(1-p_{er})^2} + 1} = \sqrt{\frac{(1-p^2)}{(1-p)^3}} \quad (17)$$

A more useful guideline for QoS measure is simply the number of transmission attempts required to reduce the probability of a packet error below a given threshold. Based upon (8), it is straight-forward to see that the number of transmission attempts N_{att} required to reduce the packet error probability below a limit given by Λ_{er} is given by

$$N_{att} \geq \frac{\log_e\left[(1-p_{er})\Lambda_{er}\right]}{\log_e(p_{er})} \quad (18)$$

Taking this result one step further, the number of attempts required to keep the overall bit error rate below a limit of Λ_{ber} in spite of unacknowledged data packets (i.e., ARQ failure) is

$$N_{att} \geq \frac{\log_e\left[(1-p_{er})\frac{\Lambda_{ber}}{p}\right]}{\log_e(p_{er})} \quad (19)$$

where p_{er} must of course be based upon the BER and the proper values for L_{crc} and L_{data}. This result assumes that the data packet is passed through even if in error if all of the ARQ attempts are exhausted. Most ARQ systems will drop the packet in error if the maximum number of re-transmission attempts is exhausted without error-free delivery. In this scenario, the probability that a packet is dropped given N_{att} attempts to deliver the packet is simply given by

$$P_{Cell\_Drop} = \left(p_{per}\right)^{N_{att}} \quad (20)$$

As clearly shown in Figure 15, the throughput efficiency plummets for a given packet-size if the BER is not sufficiently low. As suggested here, if the packet length is chosen to be the 1000 byte size called out in the IEEE802.11a standard corresponding to 8000 bits, the BER of the system must be somewhere between 10^{-5} and 10^{-6} in order to incur less than 10% throughput inefficiency due to the accompanying high PER. Although wireline and even lower-throughput wireless systems can normally achieve this irreducible BER floor, radios operating in the 16-QAM rate ¾ mode and higher (particularly 64-QAM) have an increasingly difficult chore in achieving these error floors over the multipath channel along with radio hardware imperfections. More will be said on this topic momentarily.

Another very important aspect of the underlying system PER pertains again to the QoS performance as dictated by (19). Assume for example that an acceptable system PER level for good MPEG2 video quality[8] is 10^{-7}. If the (synchronous) MAC frame length is assumed to be T_{MF}= 1.5 msec, the maximum possible time jitter that the ARQ mechanism can introduce into the video stream is $(N_{att}-1)T_{MF}$. Again for discussion purposes, assume that the maximum time-jitter[9] that the system can endure is 9 msec which translates into N_{att} = 5 maximum. For this example from (20), the PER must be smaller than 4% which translates into a maximum BER of approximately $3x10^{-6}$ for a packet size of 1000 bytes. This BER level is within the range mentioned earlier. This is very challenging to accomplish with 64-QAM over real radios as will now be considered.

Key Point: *Good system QoS performance requires that the complete system (e.g., PHY, MAC, radio) be designed with an error-floor at sensitivity that is commensurate with tolerable PER and packet length selection.*

[8] Different video decoders handle bit errors differently! Unless "perfect" packet delivery can be achieved, the system BER performance must be sized with the specific MPEG2 decoder in mind.

[9] This may be limited by actual time-jitter specifications on the application in question or by the buffer size that can be monetarily afforded.

3.1.1 Irreducible Error Floors in Real-World Radios

Pressed sufficiently hard, any real-world radio exhibits an irreducible error floor. As alluded to in the previous section, this error floor must be sufficiently low in the context of the other key system parameters in order to achieve the end system objectives.

Figure 16 Viterbi Output BER for Otherwise Ideal IEEE802.11a with Residual Phase Noise Present 64-QAM 3/4 Mode [10]

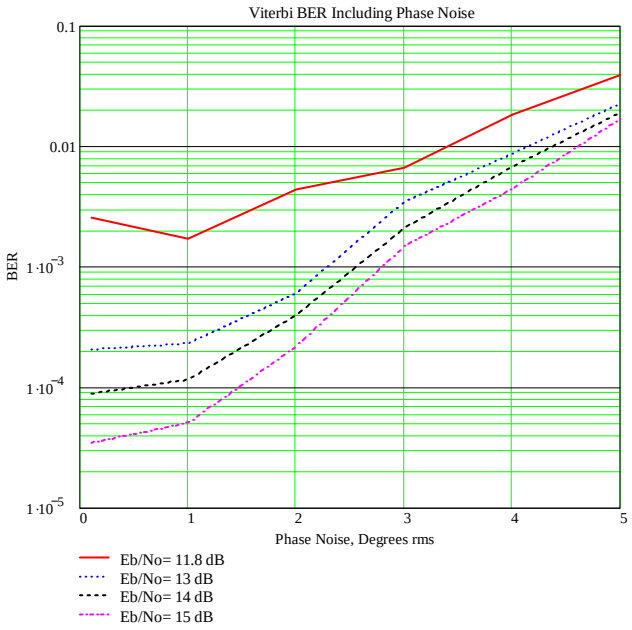

The height of the mountain is particularly high in the context of the 64-QAM operational modes. As evidenced in Figure 16, achieving a BER of 10^{-5} to 10^{-6} is very challenging, particularly when these results still do not include the additional degradations from the channel multipath and nonlinearities in the transmitter and receiver. The error floor situation is less severe in the 16-QAM ¾ case as shown in Figure 17, but even there, the BER goes up very quickly with increasing phase noise even for the E_b/N_o= 10 dB curve.

The closing point worth making here is that the Air5 system has been designed to operate at a sensitivity level corresponding to a BER between 10^{-3} and 10^{-4} even for the high-QoS mode requirements operating with 64-QAM rate ¾ . This translates into far less strenuous requirements on the radio performance and consequently lower cost and power consumption for the radio electronics. The Air5 system routinely operates at PER levels typically below 1% even over heavy multipath channels, and this is in part a key ingredient of Air5's excellent QoS performance. This perspective also provides insight into why most if not almost all IEEE802.11a systems operate at throughput rates corresponding to 16-QAM rate ¾ and below when they are required to guarantee "reasonable" QoS-like performance even on dedicated links.

Figure 17 Viterbi Output BER for Otherwise Ideal IEEE802.11a with Residual Phase Noise Present 16-QAM 3/4 Mode

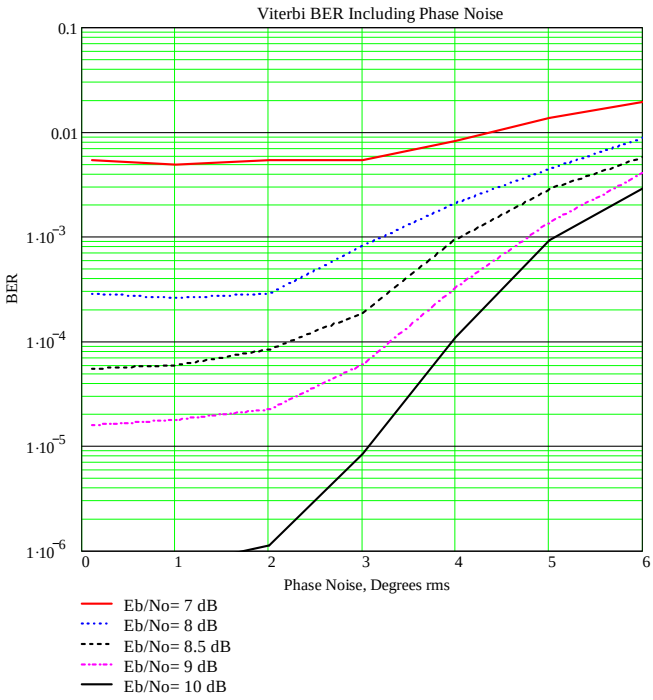

Key Point: *The Air5 system has been designed to operate at a sensitivity level corresponding to a BER between 10^{-3} and 10^{-4} even for the high-QoS mode requirements operating with 64-QAM rate ¾ . This translates into far less strenuous requirements on the radio performance and consequently lower cost and power consumption for the radio electronics. The Air5 system routinely operates at PER levels typically below 1% even over heavy multipath channels, and this is in part a key ingredient of Air5's excellent QoS performance.*

[10] Internal memo M13526

4. General Diversity Benefits in Severe Multipath Channels

Measurements like those made in Section 2 revealed early on that indoor communication throughput and reliability was limited more by multipath factors than by signal-strength issues. While this was not always the case, if severe multipath was present, even short-range communication links were always compromised severely.

In large areas where multipath is particularly poor, the propagation loss can be considerably less than predicted by free-space predictions. Also as discussed in Section 2, propagation losses through most common building materials are almost the same at 5 GHz as they are at 2.4 GHz. On the other hand, signal reflections from most materials are worse at 5 GHz than at 2.4 GHz as supported by Table 1. For systems that exploit multipath like Air5, this is a significant propagation benefit whereas for other systems that do not exploit multipath, this represents a further performance degradation.

One of the first investigations carried on at Magis was to develop a thorough understanding of what communication theory limits applied to the indoor wireless channel, and given that understanding then determine what cost-complexity solutions could be designed and built to achieve the end objectives. To this end, the Magis team built a multi-antenna array signal capturing system to assist in this endeavor.

The IEEE802.11a standard utilizes a standard k=7 Viterbi convolutional encoder that primarily encodes across the 48 data-bearing OFDM frequency bins. Adequately severe frequency-selective fading will of course break the code thereby causing burst-errors to occur. In order to assess the potential viability of different solutions, a cutoff-rate based scoring metric was used in conjunction with the antenna array signal capturing system thereby permitting the assessments to be done without a fully operational IEEE802.11a system in hand. This was an important step at the time because no IEEE802.11a systems or appropriately-equipped test equipment even existed when Magis began its development activities.

The channel cutoff rate is shown versus E_b / N_o for several square QAM signal constellations in Figure 18 as given by (21) for square-QAM signal constellations.

$$R_o(N_o) = -Log_2\left[\frac{1}{M^2}\sum_{ki}\sum_{kj}\exp\left(-\frac{\left\|S_{ki}-S_{kj}\right\|^2}{4N_o}\right)\right] \qquad (21)$$

This relationship for R_o can be summed over all of the OFDM bins in the modulation bandwidth including the effects of frequency-selective fading in order to estimate the impact of a given fading characteristic upon the information capacity of the impaired channel. This approach is helpful in that it decouples the limitations imposed by the channel from any specific algorithmic implementation.

Figure 18 Channel Cutoff Rate R_o for Several Square QAM Signal Constellations in AWGN

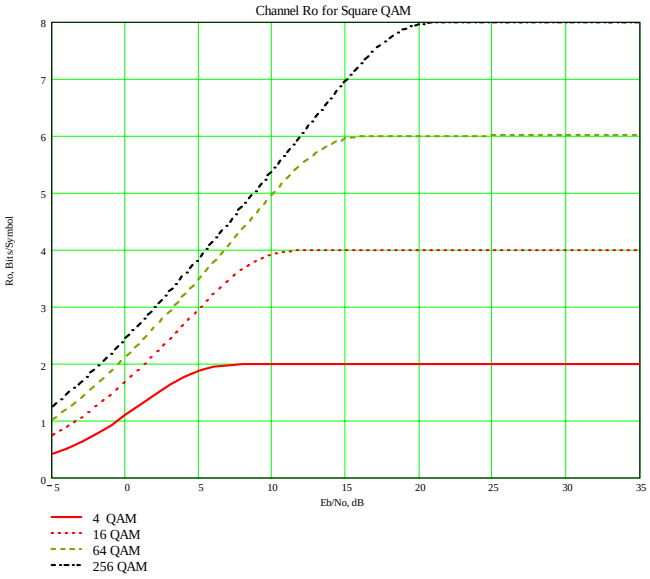

One channel sweep in which the signal strength versus frequency for 4 elements of the antenna array are simultaneously plotted is shown in Figure 19 through Figure 21. The overall channel R_o was then calculated based upon assuming different maximum C/N_o levels across the modulation bandwidth. High-level assessments like this allowed the Magis system team to determine plausible engineering solution paths (i) while using the actual real propagation channel (ii) prior to expending the extraordinary effort required to build a complete hardware prototype of our Air5 system. The fact that the Magis system development was based upon real-world actual 5 GHz wireless channel behavior rather than published static channel models cannot be over-

stressed. The behavior of the 5 GHz wireless channel is so complicated that the Magis development team could have easily been totally consumed with just the effort required to model the channel; given that the 5 GHz channel is what it is and it is immediately available and in one's midst, we took the much more simple approach of using the channel directly during the development work.

Figure 19 Baseline Multi-Element Array Channel Sweep (Only 4 Elements Shown)

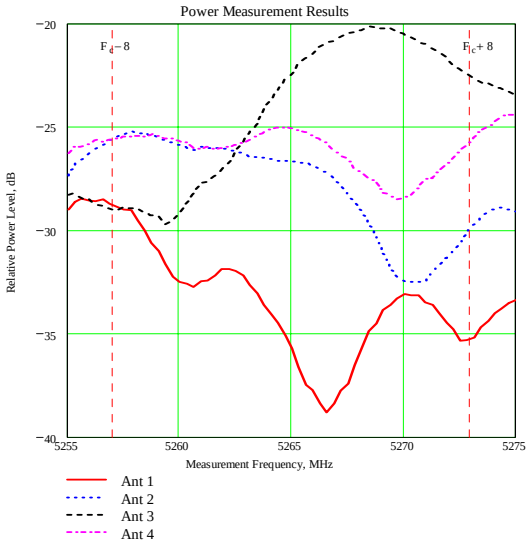

Figure 20 Computed R_o Using the Baseline Sweep from Figure 19 Assuming C/N_o= 20 dB Maximum.

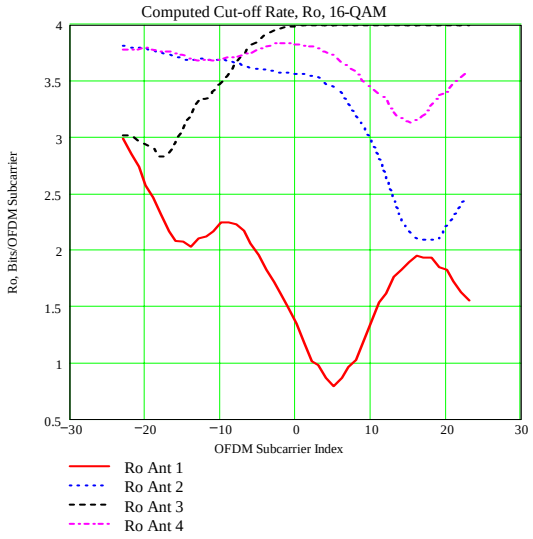

Figure 21 Computed R_o Using the Baseline Sweep from Figure 19 Assuming C/N_o= 10 dB Maximum.

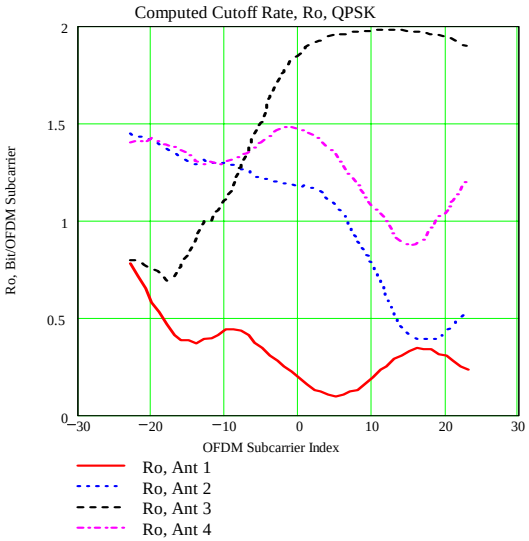

Key Point: *Single antenna receive systems that only provide one "look" at an incident signal wavefront potentially suffer a substantial degradation in sustainable information throughput due to frequency-selective fading.*

4.1 Key to Improved Performance with Multipath: Increased Dimensionality

As supported by the preceding discussion and in Section 2.2.3, a single-antenna system is generally not sufficient for dealing with the potentially severe multipath that arises over 5 GHz (and quite frankly 2.4 GHz channels as well) indoor wireless channels when attempting to communicate at the maximum throughput rates dictated in the IEEE802.11a standard plus QoS. Whether the solution lies in the traditional terminology of diversity (spatial, frequency, polarization, time, code, etc.) or the more recent vernacular of the day (array processing, space-time coding, etc.), all of these techniques attempt to increase signal processing dimensionality in order to obtain additional information about the signal before making data decisions.

In principle, these same techniques could be implemented for 2.4 GHz systems also, provided that (i) a synchronous MAC layer is available like in Air5,

and (ii) additional enhancements are added beyond those available in IEEE802.11a/g.

Estimation theory clearly teaches that exploitation of additional dimensionality in a system (when present) always leads to better system performance. The question remains as to how much additional dimensionality is available and needed for the 5 GHz channel. Consideration of this question is presented in Section 7. Aside from some very basic assumptions, the conclusion reached is that a minimum of 4 receive antenna elements are needed in order to deliver acceptable channel reliability with the multipath model assumed. Ideally, each antenna can provide a statistically independent look at the incident 5 GHz signal wavefront. Even if the signals received at each antenna element have some correlation, many other studies have shown that the diversity benefits are still very significant.

5. Air5 Enhancements Beyond IEEE802.11a

Bits and pieces of the Air5 story have been alluded to in earlier sections. The underlying concepts will be more closely identified within the Air5 context in the material that follows.

5.1 Synchronous Medium Access Control

A synchronous MAC is crucial to QoS. CSMA/CA MAC protocols like IEEE802.11 are fundamentally unable to guarantee QoS unless custom alterations are made to the long-standing protocol and only one stream is being transported on the network.

A synchronous MAC also delivers many other side benefits to the overall network performance. Some of these benefits include:

- Fundamentally better data payload efficiencies are delivered and consequently better spectrum utilization. This is a particularly serious issue when more than one user is trying to use a single network;
- Since usage of the entire wireless network is fully orchestrated by the Access Point (AP), many performance benefits in the PHY layer can be exploited as described further in Section 5.2.1.
- MAC synchronicity can be exploited in the PHY layer algorithms to improve data throughput rate and range performance.

Additional comments are provided in the sections that follow.

5.1.1 MAC Efficiency Versus Packet Length

Data packet length and its role in QoS, throughput efficiency and BER was developed at length in Section 3.1. As discussed there, IEEE802.11-based MACs are forced to use data packet lengths on the order of 1000 to 1500 bytes in order to deliver payload throughput efficiencies on the order of 60% for a single-user network. This has serious repercussions when it comes to the residual error floor of the wireless system at the higher signaling rates.

In sharp contrast, Air5 utilizes a data packet size more on the order of 200 bytes that allows the radio system to operate at a residual error floor that is about 2 orders of magnitude worse than that permitted for an IEEE802.11-based system (translating into more range and far better QoS). The Air5 system still delivers high throughput efficiency even with the much smaller packet size because the Air5 MAC is synchronous.

5.1.2 Channel Coherence Versus MAC Frame Length

Channel coherence factors were mentioned in Section 2.2. The Air5 system has been designed to accommodate motion within the propagation volume due to people moving, doors closing, etc. The synchronous MAC combined with PHY-layer enhancements permits the Air5 system to continually estimate and anticipate dynamic changes in the channel far beyond that possible with IEEE802.11-based systems.

5.2 Physical Layer Enhancements

The physical layer includes the radio electronics as well as all of the digital signal processing associated with the modem functions. Some comments have already been offered earlier regarding phase noise performance and synchronous versus asynchronous MAC operation. This section primarily addresses the baseband digital signal processing aspects of the system.

5.2.1 Estimation Theory Factors

The nonsynchronous substantially uncoordinated network attributes of IEEE802.11-based systems impose many difficulties on high-throughput reception at 5 GHz. As alluded to earlier, the Air5 system typically operates at a PER < 1% even for the 64-QAM rate ¾ mode, in no small part a benefit derived from the synchronicity of the Air5 MAC. Once a Remote Terminal (RT) has joined an Air5 network, the signal parameter uncertainties that must be dealt with by any Air5 terminal are much less demanding than those typical of an IEEE802.11-based system as partially summarized in Table 3. These main signal characteristics are briefly discussed in greater detail following.

Table 3 Comparison Between IEEE802.11 and Air5 Signal Parameter Uncertainties

Signal Attribute	IEEE802.11	Air5
Signal Time of Arrival Uncertainty	8 µsec to ??	± 100 nsec
Signal Amplitude Uncertainty at Receiver	Unknown	±2 dB
Frequency Uncertainty at Receiver	up to ± 40 ppm	± 0.01 ppm
Channel Characteristics	Unknown	Tracked (Proprietary)

5.2.1.1 Signal Time of Arrival

Once an RT has joined an Air5 network, all of the RF channel activities and resources are choreographed by the AP. An RT's use of the wireless channel whether for reception or transmission is scheduled in advance by the AP's MAC and there is very little uncertainty as to when signal activity will occur. In contrast, IEEE802.11-based networks utilize time gaps (e.g., DIFS) to arbitrate use of the wireless channel resources and there may be large time gaps if network activity is suspended for any length of time. As a result, there are large uncertainties associated with the time of arrival of a signal at the receiver. The probability for false detection of a signal increases with time as well.

In contrast, the Air5 MAC orchestrates user time slots down to a precision of a few microseconds for down-link as well as up-link communications. The system benefit is less overhead associated with dead-time over the channel as well as much lower receiver false-detection rate (in part also due to the fact that each receiver is presented a signal level that it expects due to the closed-loop power control that is built into the system.)

Signal time-of-arrival is orchestrated by the Air5 MAC to a precision within a few microseconds.

5.2.1.2 Receive Signal Amplitude

In an IEEE802.11-based network, the signal level received in any given time-slot is unknown a-priori. The received signal level may be as low as -95 dBm or as strong as -20 dBm. This poses extreme difficulty for any practical radio receiver because the instantaneous dynamic range is at best[11] approximately 35 dB. If the receiver AGC gain is too high, a strong signal will be driven into clipping whereas if the AGC gain is too low, a weak signal will be lost in the noise floor of the receiver.

In the case of 64-QAM rate ¾, a conservative figure for the required SNR is 25 dB. If the receiver is at best capable of delivering an SNR of 35 dB, the system must select its AGC setting within ±5 dB in order to avoid performance degradation. In the case of IEEE802.11-based systems, the preamble detection, signal level estimate, and coarse frequency estimate must all be done within the first few short-symbols which is fairly challenging. With the normal frequency-selective fading that is present for the indoor channel, the system's ability to counter the

[11] Common limiting factors include local oscillator phase noise, nonlinearities at the transmitter as well as the receiver, channel delay spread, quantization noise, etc.

selective fading is reduced almost dB-for-dB if AGC gain setting errors are present.

The Air5 MAC orchestrates the time and signal-level scheduling for every time-slot in advance thereby making it common place for received signal levels to be within roughly ±2 dB of that expected by any Air5 receiver.

5.2.1.3 Frequency Error

The IEEE802.11a standard permits frequency errors as large as ±20 ppm to occur at each end of the wireless link. In total, this means that the frequency error could be as large as ± 214 kHz (at 5.35 GHz). IEEE802.11-based systems perform a coarse frequency estimate during the short-symbol portion of the signal preamble and must effectively remove most of the error prior to the long-symbol portion of the preamble in order to obtain a quality estimate of the channel using the long-symbol region.

The Cramer-Rao bound combined with the achievable hardware performance limits the achievable frequency estimation accuracy. In order to get a crude estimate of the allowable frequency errors involved for 64-QAM operation, assume that a phase difference across the 16 µsec long-symbol region must be kept to less than[12] 4 degrees. This equates to an allowable frequency error of at most ± 694 Hz.

The Cramer-Rao (CR) bound for the frequency estimation error when amplitude and phase are also unknown is given by

$$\sigma_f = \frac{1}{2\pi T_m}\sqrt{\frac{1}{\frac{E_s}{N_o}N_c}} \qquad (22)$$

where E_s is the symbol energy per subcarrier, N_o is the AWGN power spectral density, and N_c is the number of subcarriers involved (52). The quantity predicted by (22) has units of Hz rms. In order to have negligible performance impact on the system due to residual frequency error, the quantity predicted by (22) should be a factor of 2 or 3 better than the requirement (in this case, ± 694 Hz for 64-QAM).

For IEEE802.11a, the measurement time is at most 6.4 µsec for the T1 and T2 long-symbol portion of the preamble, but it may be taken to be 3.2 µsec due to other implementation considerations. The CR bound is shown versus SNR in Figure 22.

Based upon Figure 22, a factor of 2 translates to a minimum receive SNR of better than 25 dB. If the frequency estimate can be made over the full 6.4 µs available, the theoretical requirement on SNR is reduced by 3 dB.

Figure 22 Cramer-Rao Bound for Frequency Estimation on Preamble ($T_{measure}$= 3.2 µsec)[13]

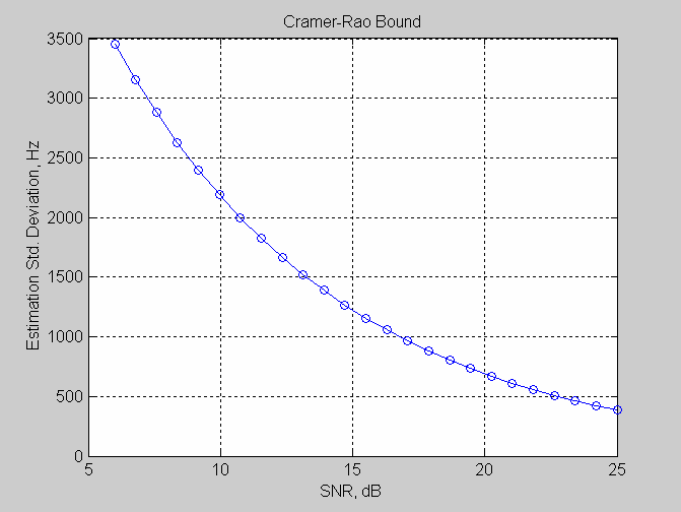

Magis defines the system sensitivity floor to be at a BER of 10^{-4} after the Viterbi decoder. With a standard IEEE802.11a implementation, this floor should be achieved at an SNR of approximately 22.5 dB for 64-QAM R= ¾ at which point the CR-bound for the frequency estimation error is 517 Hz rms from Figure 22. For 16-QAM operation at R= ½, the CR-bound on the error is 1.71 kHz rms (SNR approximately 12.1 dB assumed) again from Figure 22. These frequency errors must somehow be eliminated or dealt with in the OFDM receiver or the channel estimate made during the long-symbol portion of the preamble will become obsolete very quickly.

[12] For a fixed frequency error, the "phase ramp" is averaged over the long-symbol region thereby resulting in a net error of one-half the peak-to-peak phase difference involved.

[13] M13459 PlotCR.m

Losses Associated with Initial Frequency Error

The open literature most often points to coherence loss as one of the major issues when an initial frequency error is present in an OFDM system. Since coherent demodulation of each OFDM subcarrier is also involved, constellation point skew and rotation is also another major issue that is actually more dominant. These impairments are considered separately in the following sections.

Orthogonality Related Losses- Loss of Bin Energy

Due to the loss of orthogonality associated with an initial frequency error, some loss arises due to effectively a matched-filter mismatch[14]. For reasonable frequency errors, this is a very small loss as given by

$$Loss(\delta F) = -20Log_{10}\left[\left|\frac{\sin(\pi T_s\ \delta F)}{\pi T_s\ \delta F}\right|\right] \qquad (23)$$

With T_s= 3.2 µsec and δF= 30 kHz for example, the loss is still only 0.132 dB and is therefore completely negligible.

Orthogonality Related Losses- Subcarrier Cross-Coupling

As developed elsewhere[15], the interference level from an OFDM bin separated by w bins from the bin of interest when a frequency error of δF is present is given by

$$CIR(w) = -20Log_{10}\left\{\left|\frac{\sin\left[\pi T_s\left(wF_B + \delta F\right)\right]}{\left[\pi T_s\left(wF_B + \delta F\right)\right]}\right|\right\} \qquad (24)$$

where F_B is the OFDM subcarrier spacing (312.5 kHz) and T_s is (F_B)$^{-1}$. The interference from the adjacent OFDM subcarriers will be statistically independent and the accumulation of interfering terms will be most pronounced for the center-bin location. The total carrier-to-interference ratio (CIR) due to the loss of orthogonality is shown graphically in Figure 23 for the IEEE802.11a waveform.

For 64-QAM rate ¾ at an SNR of approximately 25 dB at sensitivity, ideally the interference caused by frequency-error related orthogonality loss is 10 dB lower or in this case 35 dB. Based upon Figure 23 the frequency error should be kept to less than roughly 3 kHz to deliver this level of performance. Compared to the CR-bound of 517 Hz at this SNR, this is "5.8-sigma" which means that so long as the frequency estimator is efficient, performance for the 64-QAM rate ¾ mode as far as initial frequency error is concerned will be very adequate.

For 16-QAM rate ¾ at an SNR of approximately 14 dB at sensitivity, ideally the interference caused by frequency-error related orthogonality loss is 10 dB lower or in this case 24 dB down. Based upon Figure 23 the frequency error should be kept to less than roughly 10 kHz to deliver this level of performance. Compared to the CR-bound of approximately 1200 Hz from Figure 22, this is "8.3-sigma" which means that so long as the frequency estimator is efficient, performance for the 16-QAM rate ¾ mode as far as initial frequency error is concerned will be very adequate.

Figure 23 CIR Impact on Center OFDM Subcarrier Bin due to Loss of Orthogonality from Initial Frequency Error[16]

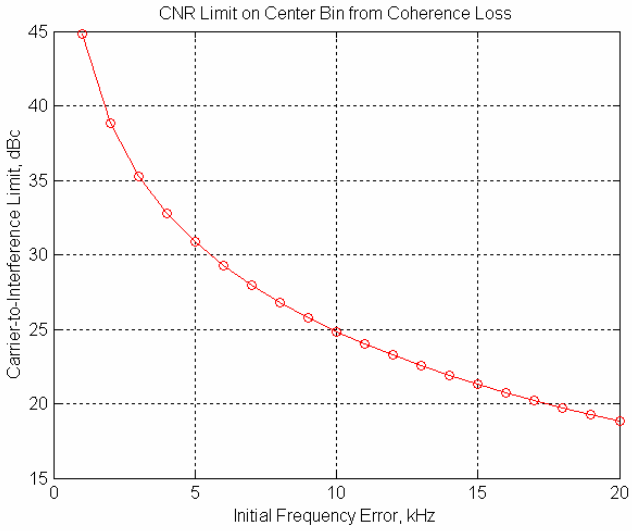

[14] Internal memo, M13415
[15] Internal memo, M13415
[16] M13458 OrthoLoss.m

Multi-Variate Uncertainties (Fisher Information Matrix)

When more than one signal parameter is in question, the Cramer-Rao bound is rigorously defined in terms of the Fisher information matrix [6,7]. Although of mathematical interest, the system performance will be more likely determined by the finite instantaneous dynamic range of the receiver and other implementation issues. This more rigorous pursuit of the estimation problem has therefore not been presented here.

Local Oscillator Phase Noise Limitations on Channel Estimate

Local oscillators (LO) within the radio portion of the system have non-ideal phase noise characteristics that affect system performance. The effects of phase noise on BER was looked at earlier in Section 3.1.1. LO phase noise content also affects the frequency estimation accuracy and channel estimation accuracy possible from observing the long-symbol portion of the IEEE802.11a preamble.

Assuming that the composite SSB phase noise spectrum of the system is given by L(f) (including both transmit and receive ends of the system), and a finite observation time of T_m= 6.4 μsec, it can be shown that the variance of the LO phase that is observed over the channel estimation time is given by

$$\sigma^2_{\Delta\theta} = E\left\{\left[\int_0^{T_m} \theta(t+\tau)\,d\tau\right]^2\right\} \qquad (25)$$

where the assumption is that no adjacent OFDM subcarrier tones are present. In a normal real-world case, the presence of adjacent OFDM subcarriers results in reciprocal mixing with the LO phase noise spectrum also thereby contributing additional interference. Assuming that all of the OFDM subcarriers contain the same average signal energy, the signal variance due to the LO phase noise for the k^{th} subcarrier frequency bin is given by[17]

$$\sigma^2_{Tot}(k) = \int_{-\infty}^{+\infty} L(f)\left[1-\frac{\sin^2(\pi f\, T_m)}{(\pi f\, T_m)^2}\right] + \sum_{\substack{n \\ n\neq k}} \int_{-\infty}^{+\infty} L(f)\left[\frac{\sin(\pi f\, T_m)}{\pi(f+(l-n)\Delta F)}\right]^2 df \qquad (26)$$

where ΔF = 312.5 kHz.

As shown by the first integral portion of (26), the finite observation time T_m introduces a highpass filter characteristic into the integrand which helps substantially in reducing the impact of non-ideal phase noise. The second integral portion of (26) is the reciprocal mixing contribution from all of the other OFDM subcarriers effectively mixing with the phase noise spectrum sidelobes.

Hypothetical or actual phase noise power spectral densities L(f) can be used in (26) to quickly assess how good the LO phase noise spectrum needs to be for a given overall system budget.

5.2.2 Multiple-Antenna Signal Processing

The spatial wavefront signal processing[18] technology is probably the single most important aspect of Air5 that allows it to deliver reliable high-throughput wireless communication unlike any thing else currently available. In its most simple terms, if a single antenna system has no appreciable signal energy incident at its one antenna due to multipath, no amount of fancy signal processing will pull the bits out of the noise.

Earlier sections of this memorandum have presented the basic need for multi-element receive signal processing from several perspectives. The results presented in this section are based upon the more rigorous computational assessment of the problem provided in Section 7.

The approach taken in Section 7 assumes a completely random incident wavefront composed of 3 multipath rays on a linear array of antenna elements. In order to separate out the theoretical limits of the problem from the complexity and implementation related factors, a channel cutoff-rate (R_o) approach to the problem was used. This approach allows the

[17] Internal Memo M13651

[18] Referred to as CSWR for Crawford Spatial Wavefront Receiver

problem to be considered independently of the forward-error protection used, interleaver depths, etc. and only requires that the underlying signal waveform be roughly known; in this case IEEE802.11a OFDM. No losses due to phase noise, nonlinearities, frequency errors, etc. are of course included in the assessment.

The R_o-based approach is motivated by the random coding bound which stipulates that the average probability of error can be bounded[19] as

$$\overline{P(e)} \prec 2^{-N[R_o - R_N]} \qquad (27)$$

Other metrics could of course be used in this assessment, but this approach is consistent with other work that had already been done within the company.

As argued in Section 7, a minimum of 3 antennas are required in order to deliver the needed reliability assuming "everything else" is ideal including the FEC and interleaving. With a more realistic real-world system (like the IEEE802.11a OFDM waveform), it is very straight-forward to argue that a minimum of 4 receive antennas would be required to achieve the same degree of reliability.

The Air5 CSWR architecture is fundamentally critical to delivering high communication reliability, and with that the ability to deliver unprecedented throughput rate and video support.

6. References

1. J. Musil and F. Žáček, Microwave Measurements of Complex Permittivity by Free Space Methods and Their Applications, Elsevier, New York, 1986.
2. J. Medbo, "Simple and Accurate Path Loss Modeling at 5 GHz Indoor Environments with Corridors", VTC 2000
3. Wilson, R., "E10589 Propagation Losses Through Common Building Materials", August 2002
4. J.K. Holmes, Coherent Spread Spectrum Systems, John Wiley & Sons Inc., 1982, Chapter 9
5. Villanese, F., et al., "Pedestrian-Induced Fading for Indoor Channels at 2.45, 5.7 and 62 GHz", VTC 2000
6. Y. Jiang, "On the True Cramer-Rao Lower Bound for the DA Joint Estimation of Carrier Phase and Timing Offsets", ICC 2000
7. L. Scharf, *Statistical Signal Processing*, 1991, Addison-Wesley
8. J.M. Wozencraft, I.M. Jacobs, *Principles of Communication Engineering*, 1965, John Wiley & Sons
9. J.G. Proakis, *Digital Communications*, 4th Edition, 2001, McGraw-Hill Book Co.
10. S.G. Wilson, *Digital Modulation and Coding*, 1996, Prentice-Hall

[19] [8], page 315, [9] page 396, [10] page 331

7. Appendix: Multiple-Input Receive-Side Coherent Combining

Diversity reception has long been known as a very effective technique for combating the severe frequency-selective fading that occurs over multipath channels. Given N receive antenna elements, the optimal system would require N complete receivers. In the case of OFDM, the signal processing would optimally combine the frequency-bin information from each receiver in a bin-by-bin optimal manner. For an IEEE802.11a-based system that uses 52 OFDM subcarriers, this system would continuously compute 52 x N complex weighting coefficients as part of the combining process. Such a system would however be far too complex and expensive for the applications being considered.

A reduced-complexity solution is considered within this appendix where the combining is limited to antenna-based combining thereby only involving N complex weighting coefficients. Only the uniform linear array (ULA) antenna configuration will be considered.

A single planar signal wavefront is shown incident on a ULA here in Figure 24. Signal multipath leads to additional planar signal wavefronts being incident on the ULA from different direction-of-arrivals (DOA), amplitudes and RF phases thereby causing constructive and destructive interference at each antenna element as a function of frequency. The frequency-selective fading that results at each individual antenna was mentioned earlier in Section 2.2. The incident signal amplitude and phase due to a specific incident signal wavefront is given by

$$C_{k,n}(f) = A_k \exp\left[-j\frac{2\pi}{c}(F_c + f)\cdot\left[L_k + (n-1)d\sin\left(\theta_{DOA_k}\right)\right] + j\phi_k\right] \quad (28)$$

Figure 24 Uniform Linear Array with Single Incident Planar Wavefront

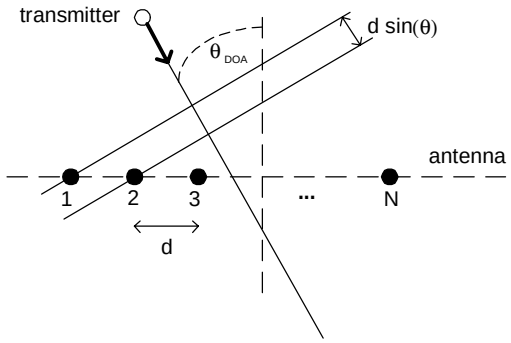

where

A_k	signal amplitude (voltage)
c	speed of light (3.0 e8 m/s)
F_c	signal center frequency (e.g., 5.25 GHz)
f	offset frequency from carrier, Hz
L_k	assumed path length, m, between transmitter and antenna #1
n	antenna index
d	antenna spacing, m
θ_{DOA_k}	angle of arrival for k[th] planar wavefront
ϕ_k	RF phase of k[th] planar wavefront at first antenna

The superposition of the different interfering signal wavefronts at the n[th] antenna leads to a combined signal amplitude of

$$H_n(f) = \sum_{k=1}^{N_{ant}} C_{k,n}(f) \quad (29)$$

and when antenna-based coherent combining is used, the net result is a channel transfer function given by

$$Hc(f) = \sum_{n=1}^{N_{ant}} \alpha_k H_k(f) \quad (30)$$

where it remains to find the complex weighting coefficients α_k such that the bit error rate over the multipath channel is minimized.

In a single-carrier system, it generally suffices to find the α_k such that signal power is maximized across the modulation bandwidth. For M-QAM OFDM however, the goodness criteria is considerably more complex. In order to correctly choose the α_k, the correct cost function must be chosen for the underlying optimization task.

For purposes of this discussion, the cost function that will be adopted is the total R_o[20] of the composite channel Hc(f) in (30). This cost function choice is attractive because it is tied closely to the information capacity of the composite RF channel[21] and does not presume any limitations related to coding, interleaving, etc. No claims are made here to the effect that in an actual system implementation that this is the best cost function possible.

Computer simulation was used to compare the system performance as a function of (a) number of antennas used, (b) amplitude of interfering multipath signals and (c) antenna element spacing. For each trial, 1000 randomly chosen triplets of DOAs were used in order to average over any directional preferences of the ULA.

Table 4 Simulation Cases Considered (16-QAM)

Case	Antennas	A_1-A_3	d	SNR_{dB}
1	5	1, 0.4, 0.2	1.5”	11
2	4	1, 0.4, 0.2	1.5”	11
3	3	1, 0.4, 0.2	1.5”	11
4	2	1, 0.4, 0.2	1.5”	11
5	5	1, 0.4, 0.2	1.0	11
6	4	1, 0.4, 0.2	1.0	11
7	3	1, 0.4, 0.2	1.0	11
8	2	1, 0.4, 0.2	1.0	11
9	5	1, 0.4, 0.2	2.0	11
10	4	1, 0.4, 0.2	2.0	11
11	3	1, 0.4, 0.2	2.0	11
12	2	1, 0.4, 0.2	2.0	11

A convenient approximation for the R_o function that was used in the antenna weight optimization step is

$$R_o\left(SNR\right) \approx 4\left[\tanh\left(a_1 SNR + a_0\right)\right] \qquad (31)$$

where a_1= 0.052 and a_0= 0.30 for 16-QAM. This approximation becomes increasingly inaccurate at R_o values less than approximately 2.0 but since this corresponds to a region of performance that is to be avoided, the approximation errors do not pose any difficulties.

A summary of the average R_o per OFDM subcarrier and standard deviation for each simulation case is provided in Table 5.

Figure 25 Actual Ro Versus Approximate Ro for 16-QAM

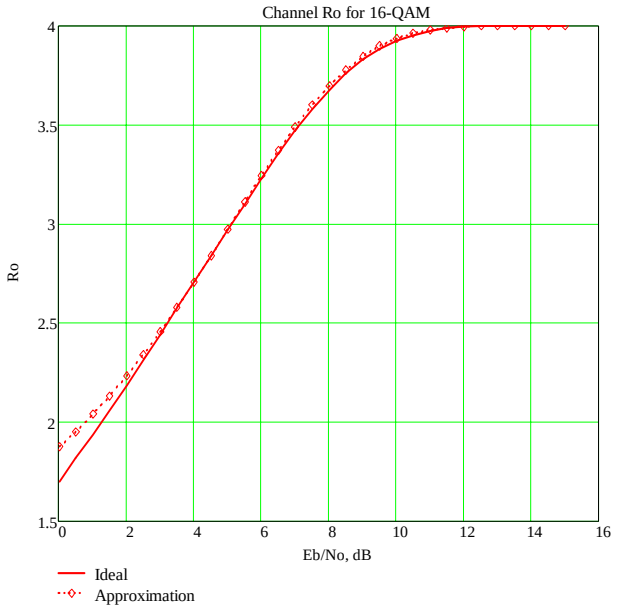

Table 5 Summary of Simulation Results

Case	Ave R_o Single	Ave R_o Combined	Std Dev R_o Single	Std Dev R_o Combined
1	**2.85**	**3.946**	**0.34**	**0.054**
2	**2.85**	**3.906**	**0.34**	**0.083**
3	**2.85**	**3.800**	**0.34**	**0.132**
4	**2.85**	**3.508**	**0.34**	**0.229**
5	**2.85**	**3.95**	**0.34**	**0.0497**
6	**2.85**	**3.911**	**0.34**	**0.0792**
7	**2.85**	**3.812**	**0.34**	**0.1251**
8	**2.847**	**3.509**	**0.34**	**0.221**
9	**2.846**	**3.949**	**0.34**	**0.051**
10	**2.85**	**3.917**	**0.34**	**0.073**
11	**2.85**	**3.816**	**0.34**	**0.118**
12	**2.85**	**3.519**	**0.34**	**0.2114**

In order to gain additional insight into the role that antenna number and spacing affect the average R_o for each scenario in Table 4, the probability density functions for each scenario were plotted in

[20] See Section 4 for some discussion of cutoff rate, R_o

[21] An assumption must be made as to what type of signal space is being used for each OFDM bin; e.g., 16-QAM

Figure 26 5 Antennas / 1.0 Inch Spacing

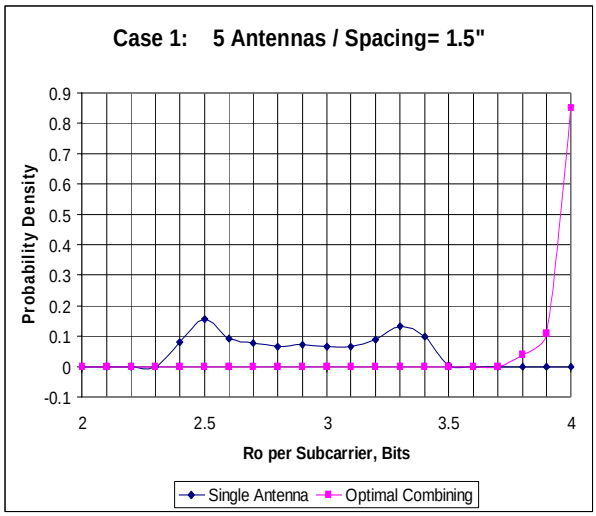

Figure 27 4 Antennas / 1.0 Inch Spacing

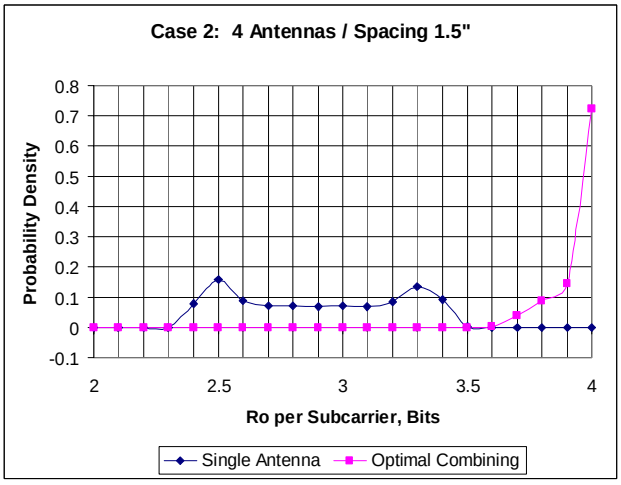

Figure 28 3 Antennas / 1.0 Inch Spacing

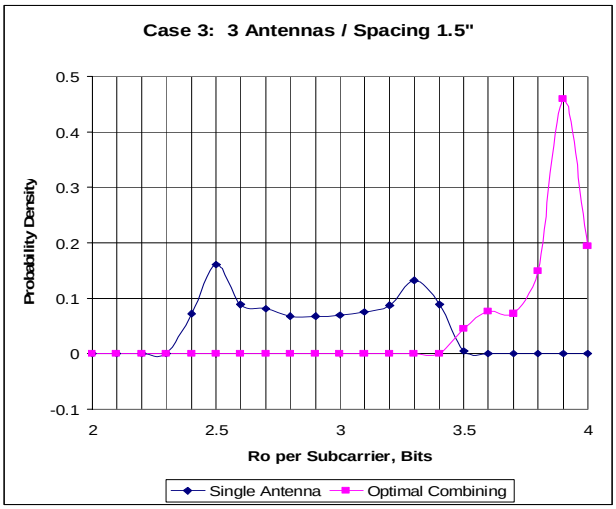

Figure 29 2 Antennas / 1.0 Inch Spacing

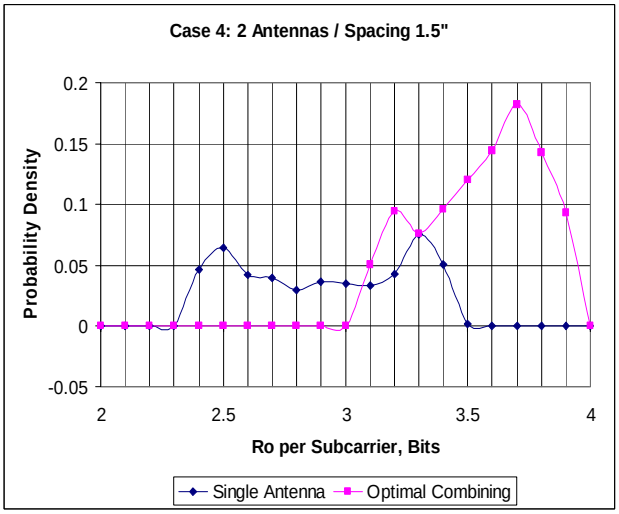

Figure 30 5 Antennas / 1.0 Inch Spacing

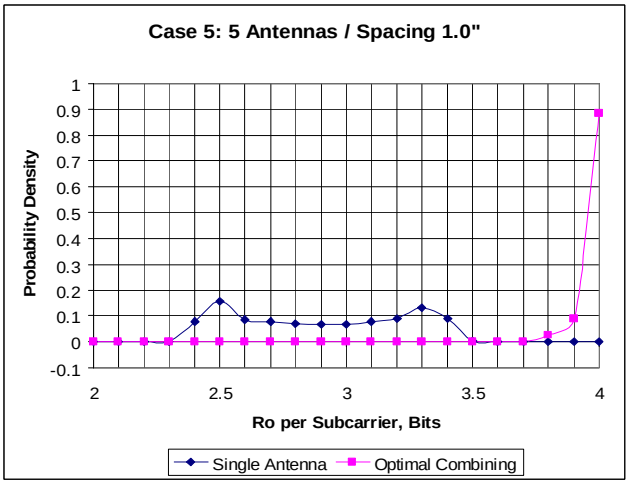

Figure 31 4 Antennas / 1.0 Inch Spacing

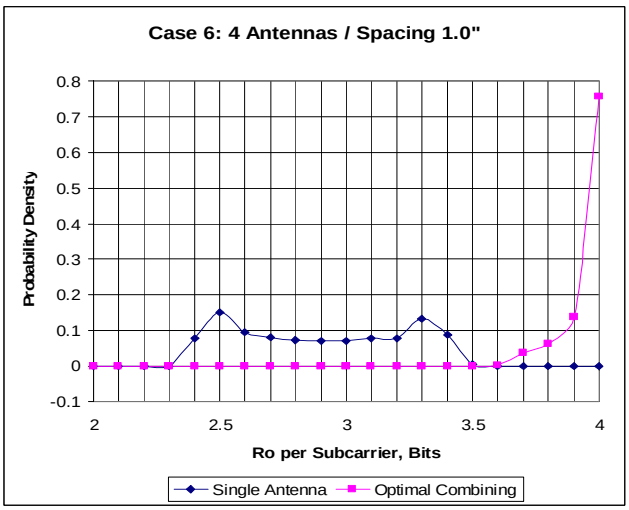

Figure 32 3 Antennas / 1.0 Inch Spacing

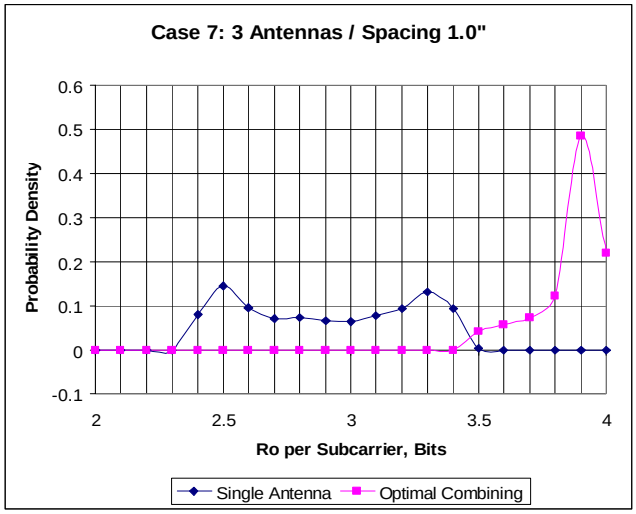

Figure 33 2 Antennas / 1.0 Inch Spacing

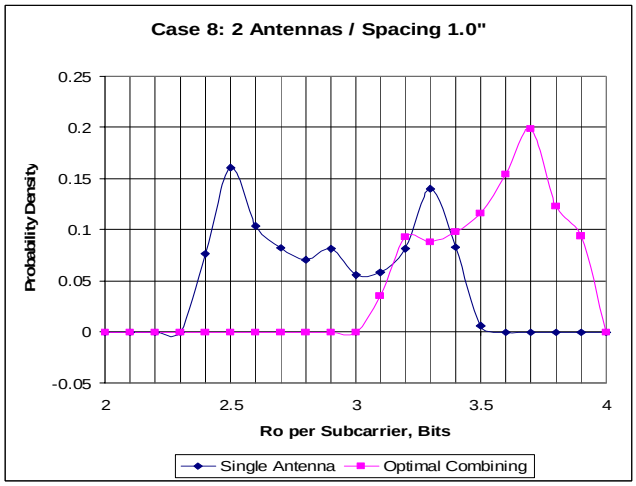

Figure 34 5 Antennas / 2 Inch Spacing

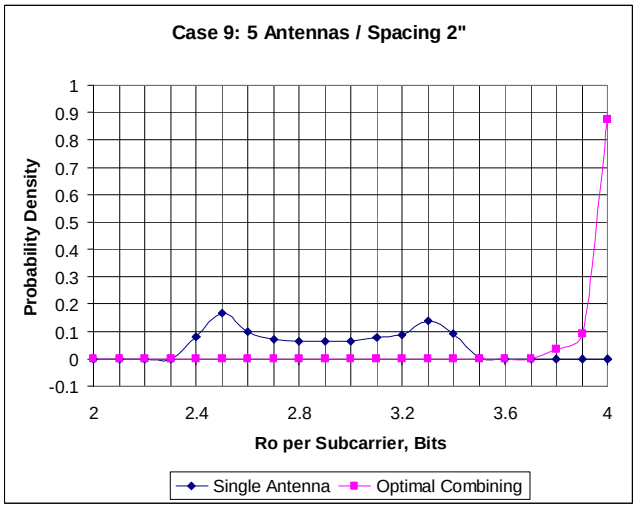

Figure 35 4 Antennas / 2 Inch Spacing

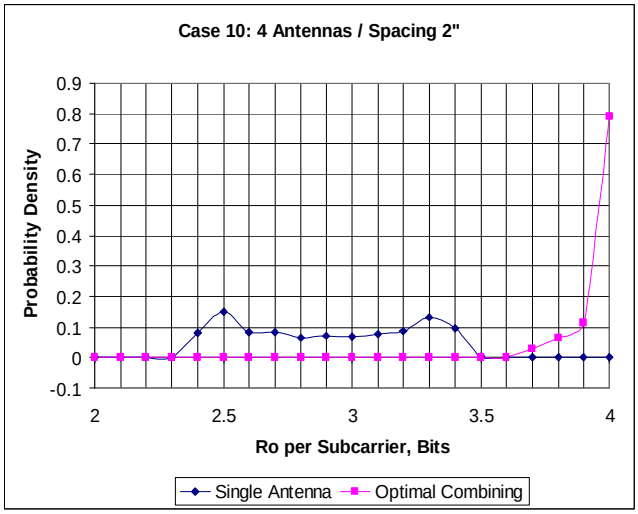

Figure 36 3 Antennas / 2 Inch Spacing

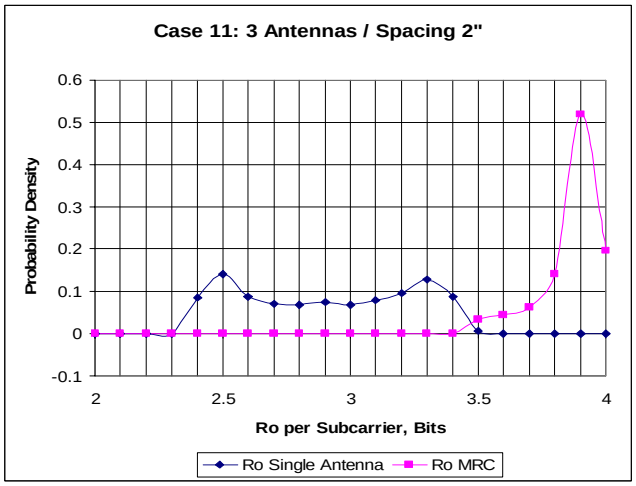

Figure 37 2 Antennas / 2 Inch Spacing

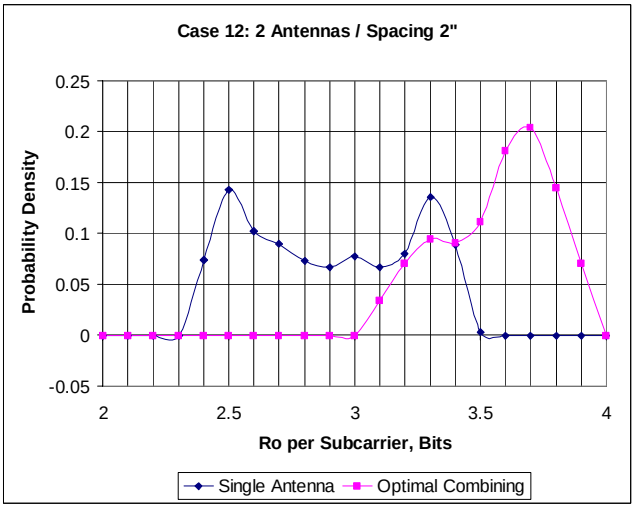

Figure 26 through Figure 37. As clearly shown in these figures, the presence of more antennas (i) increases the mean R_o for the frequency-selective multipath channel and (ii) reduces the standard deviation of R_o as well. Both of these factors contribute to a sizeable improvement in the channel reliability realized. The average R_o and standard deviation of R_o are plotted versus antenna number in Figure 38 and Figure 39 respectively.

Figure 38 Average R_o per Subcarrier versus Antenna Number and Antenna Spacing

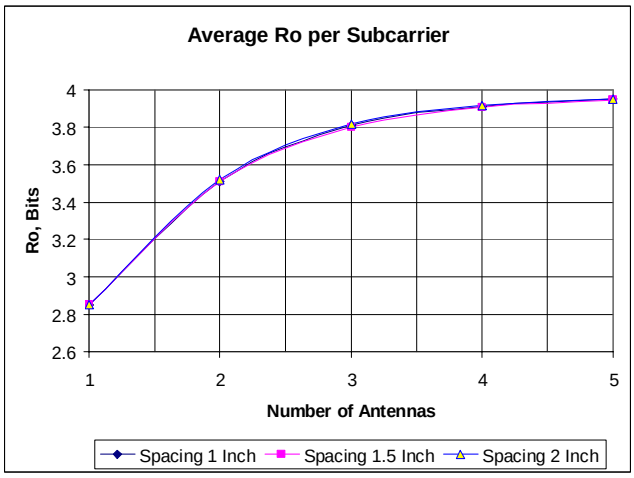

Figure 39 R_o per Subcarrier Standard Deviation versus Antenna Number and Antenna Spacing

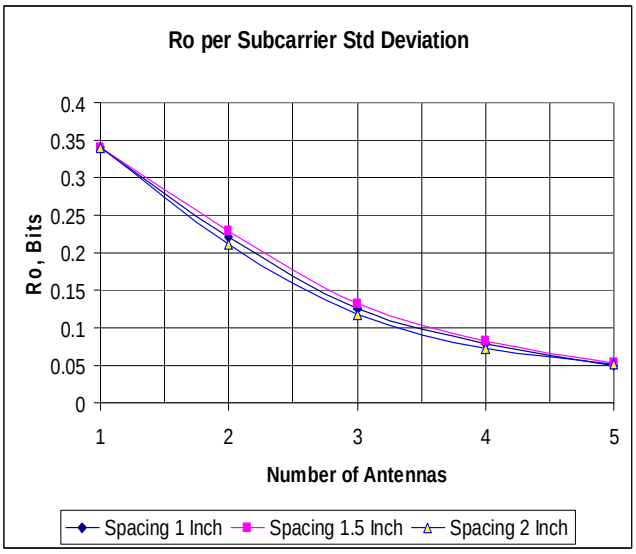

If we think in terms of the rate ¾ code that is used for the maximum throughput rate in 16-QAM and 64-QAM modes, it obviously mandatory that the average R_o be greater than at least 0.75 x 4= 3 bits per subcarrier in order to have any possibility of supporting this coding rate. Based upon Figure 38, it is then mandatory that the system use a minimum of 2 antenna elements in its array. Practical code-implementation losses combined with the standard deviation of R_o shown in Figure 39 can be used to argue conclusively that a minimum of 3 antennas is most likely required even with optimal antenna combining.

Another important point to draw from Figure 38 and Figure 39 is that the mean and standard deviation measures do not appear to be a function of the antenna spacing. This result runs counter to a number of field testing results conducted by the Magis staff and most likely arises from the exclusion of additional multipath factors that are not included in the simplistic planar wave model adopted for this analysis.

Comparison at 2.4 GHz

The same analysis was done at 2.4 GHz assuming only 3 antennas since fewer antennas per unit length are possible due to the increased wavelength as compared to 5 GHz. The R_o probability density for this case is shown in Figure 40. As evidenced here, the R_o behavior for this case is quite similar to the 3-antenna case at 5 GHz shown in Figure 28, Figure 32, and Figure 36 which is primarily due to the simplicity of the multipath model that was chosen for this analysis. Even so, it is possible to argue that the same degree of improvement is possible in the 2.4 GHz band so long as the additional physical size of the antennas can be accommodated as compared to 5 GHz.

Figure 40 2.4 GHz / 3 Antennas / 2.5" Spacing

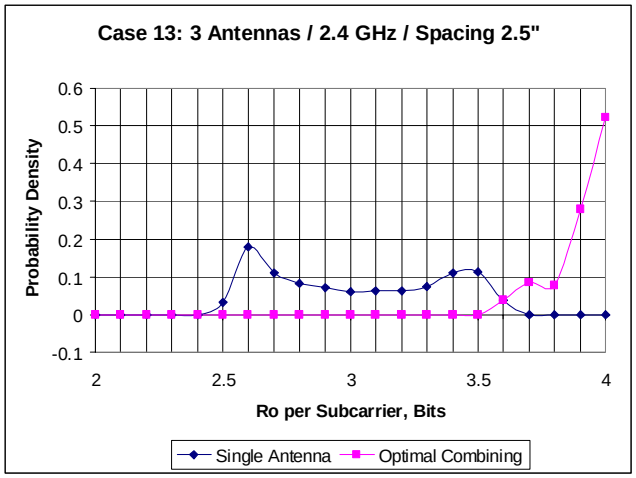

Conclusions

Depending upon the relationship between R_o and the underlying PER, the results provided here argue very strongly in favor of a minimum of 3 receive antenna elements if near-ideal antenna combining can be implemented. Other modeling factors plus the difficulty in implementing truly optimal combining will dictate the use of 4 antennas or more in any practical hardware implementation.