Multimedia Over 5 GHz Wireless Home Networks

James A. Crawford
Chief Technology Officer
Magis Networks, Inc.
jcrawford@magisnetworks.com

David N. Critchlow
Vice President of Engineering
Magis Networks, Inc.
dcritchlow@magisnetworks.com

INTRODUCTION

In recent years, wireless networks have gained popularity for data networking within the home and small office. The much larger market of delivering Multimedia content (high quality video like HDTV, SDTV as well as audio) has been left essentially untapped. Many semiconductor and product companies have attempted to serve this market with standards-based solutions, all of which are hampered by their data-centric roots. The channel reliability required to faithfully deliver a feature-length HDTV movie is orders of magnitude more demanding than traditional data applications for example. A new generation of devices is now available that have been designed from the start to meet the market requirements thereby leading to some breakthroughs in systems design and implementation.

In this paper, we explain the requirements of a Multimedia home network, and the key elements required to successfully implement a system to deliver this high quality content with guaranteed Quality of Service (QoS). The paper goes on to explain the importance of MAC/PHY/RF co-design to optimize performance in multipath environments, cost, and power consumption. We also present real-world testing results that show that when a 5 GHz wireless system is designed correctly, it can cover the entire home with guaranteed throughputs for delivery of Multimedia content.

What Has Opened This Market?

Leading consumer electronics companies and service providers are strategically positioning themselves for the connected home era by enabling rich entertainment experiences.

Although talked about for many years, only recently have market forces and available technology provided the potential for wireless multimedia delivery. The convergence of new technologies has helped to fuel the demand for home entertainment networks such as new entertainment models (e.g. on-demand viewing), new TV form factors (e.g. flat screens), and on-line multiplayer video games.

Multimedia versus Data (Streaming versus Live Video)

Companies have attempted to build multimedia networks based on computer communication protocols (e.g. 802.11x). The drawbacks of this approach are:

- These systems were designed for connecting computers only
- Host computer or processors with operating systems must always be present
- MPEG-2 content must be encapsulated into TCP/IP packets and be passed through an asynchronous network stack
- Streaming/ “pull technologies” must be used, resulting in large buffering delays and intermittent loss of signal.

The requirements of a true multimedia networking technology take the following into account:

- Interfaces designed to connect directly to A/V equipment and PCs
- Ability to operate without a host PC or operating system
- Provisions for dual data and isochronous interfaces, capable of simultaneous operation
- Direct MPEG-2 interfaces with built-in synchronization between source and destination to remove latencies and signal loss conditions. This allows for the delivery of "live" video.

Throughput versus Physical Layer Rate

The 802.11a/g standards provide for physical layer bit rates of up to 54 Mbps. At first glance, this looks like a large improvement over the existing 802.11b standard of 11 Mbps. Unfortunately, due to limitations at the MAC layer, achievable throughputs of only 20 Mbps have been demonstrated. This is a highly inefficient communication system. This is based on a CSMA-based access scheme used in 802.11x networks. Within the constraints of this standard, large overheads are incurred due to preambles, and carrier sensing periods. In an attempt to reduce this overhead as a percentage of the data payload, longer packets are sent. The problem with the use of longer packets is that they are more prone to Packet Errors over the wireless channel and hence more retransmissions occur, which then reduces the overall throughput. Throughput efficiencies of only 40-50% are achievable in point-to-point CSMA systems.

Synchronous MAC architectures used in systems such as HiSWan, Hiperlan, and the Magis Air5 system achieve higher throughputs. As an example, the Air5 system achieves throughput efficiency rates of 80% through the use of a synchronous, schedule-based TD/TDMA MAC. In this system, intra-frame carrier sense is not required, and shorter packet sizes can be used. This removes the PER problem mentioned above, and yields high throughputs with the added benefit of very low latency (on the order of a few milliseconds).

Security

With the delivery of high valued content (digital media, HDTV video, copyrighted materials), security is more important than ever. A wireless multimedia network must provide a layered security structure. An example is shown in Figure 1.

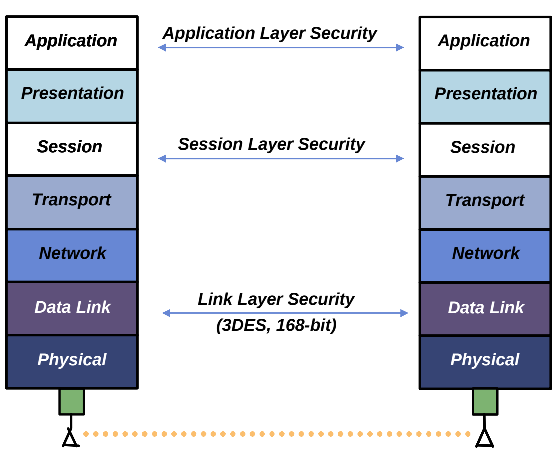

Figure 1: Air5 Security Diagram

The higher levels (Application Layer, Session Layer) must be preserved. These are all systems that exist "above" the wireless distribution system within the home. The link layer security portion should offer the following elements in order to prevent data capture and possible attacks.

- Common secret derivation and communication
- Authentication and access control
- Key derivation and periodic
- Payload encryption (3DES, AES, etc).

QUALITY OF SERVICE (QOS) – WHAT DOES IT MEAN IN WIRELESS NETWORKS?

The Importance of Latency, Jitter, and Guaranteed Delivery

Table 1 summarizes the typical QoS parameters applicable for multimedia content distribution within a home entertainment network. As the table shows, the latency and jitter requirements of these systems pose a significant design goal to any network carrying this content, especially a wireless one. A wireless network based on IP protocols will not meet all of these requirements for the reasons mentioned in the previous sections. Synchronous networks can achieve these requirements given that they are designed with the proper constraints and attention to the key latency, jitter, and PER performance parameters.

Input Parameters of Performance Testing			Output Parameters of Performance Testing			
Service	**Number of Streams**	**MAC Payload Rate (per stream)**	**Packet Size (bytes)**	**Max PER**	**Max Latency**	**Max Jitter**
Multiple Stream (1 HDTV/2SDTV)	3	35 Mbps	228	3.6* 10-5	90ms	+/-10ms
HDTV	1	19.68 Mbps	228	3.6* 10-5	90ms	+/-10ms
SDTV/DVD	1	3/8 Mbps	228	3.6* 10-5	90ms	+/-10ms
HQ Video Conf. Call	2 per call	1.5 Mbps	228	3.6* 10-5	10ms	+/-5ms
CD Quality Audio	1	256 kbps	360	5.8* 10-5	100ms	+/-10ms

Table 1 Typical QoS Parameters for Multimedia Content Distribution

Video Interfaces – Isochronous vs. Asynchronous

Transport streams (TS) are used in systems where the video is to be rendered immediately. The most prevalent examples of this are cable and satellite systems where the content is most frequently live and real-time.

The video decoder must decode at the rate at which data is being sent to it. This is the case for, among other reasons, that the content may be live and therefore must be sent and decoded at the rate at which the content is being encoded. This implies that the decoder is a "slave" to the source. Both the MPEG-2 encoder and the decoder operate at a nominal clock of 27 MHz, however these clocks are only accurate to 810 Hz. This means that the encoder's clock will drift with respect to the decoder's clock.

Typically most hardware MPEG-2 decoders contain very little buffering (i.e. memory), usually just enough for one or two frames of video. This requires the decoder to maintain fairly tight synchronization with the encoder.

There are two means for slaving the decoder to the encoder. For the purposes of this discussion they will be referred to as (i) embedded timing techniques and (ii) buffer management techniques.

Most data-centric systems (802.11x) are forced to use buffer management techniques where TS packets arriving at the decoder are placed into a buffer. The decoder reads from the buffer. Before the decoding starts, the buffer must initially fill to a watermark which once reached, the decoder will begin to pull TS packets from the buffer and begin decoding. A microprocessor or other intelligent device must monitor this buffer for underflow/overflow and make adjustments to the decoders speed or alternatively insert or drop packets.

A system design challenge arises since the buffer size depends upon the jitter introduced by the channel, the variability of the TS itself, and the jitter introduced by the interfaces between the channel and the decoder (e.g. PCI or other asynchronous data bus latencies). The buffer must be made large enough to absorb all three sources of jitter as well as distinguish the very subtle effect of clock drift between the encoder and decoder.

Existing attempts to use data-centric 802.11x protocols to deliver video, require large buffers (on the order of 32MB) which can result in delays between 2 and 8 seconds on existing products in the market today. This is an unacceptable latency, and leads to the inability to support interactive feedback or channel changing in a timely manner.

Isochronous data transport systems using embedded timing techniques have been shown to overcome the inherent system issues caused by the use of large buffers.

Certain packets within a transport stream contain a program clock reference (PCR which is a sample of the encoder's system clock. When the decoder receives a TS packet containing a PCR, it measures the number of ticks of its own system clock that have transpired since the last PCR packet. This "error term" can be used to feed a tracking circuit so as to correct the clock differences. In this way the decoder is slaved to the encoder.

In order for this process to work, the time between consecutive MPEG TS packets (referred to as inter-packet spacing) must be preserved across the channel. If the TS packets incur a jitter along the path from encoder to decoder, this jitter must be removed to within 500 nanoseconds [1].

Communication systems such as Firewire (IEEE 1394) and Magis' Air5 perform this inter-packet timing preservation. The end result is guaranteed delivery with less than 8 msec of latency, a 250 to 1000 times improvement over the data-centric approaches. Both Firewire and Air5 hardware are designed to interface directly to MPEG encoders and decoders, without the need for bus translation from the PC-world (e.g. PCI) to the A/V world (parallel or serial TS interfaces). This reduces the total bill of materials, power consumption, size, and complexity.

SYSTEM DESIGN

Understanding the Indoor Wireless Channel

The indoor wireless channel is extremely difficult to model accurately at 5 GHz because almost everything reflects incident signal energy to some degree or another. Since a quarter-wavelength at this frequency represents only about 1.5 cm in free-space, specular as well as diffuse multipath abounds in most indoor environments.

Contrary to the popular power-law based propagation model predictions that are based upon CW measurements with substantial spatial averaging, typical indoor frequency-selective fading can exhibit

- frequency nulls greater than 15 dB in depth,
- flat frequency-fades across the entire modulation bandwidth,
- fade durations that can last seconds and even minutes at a time.

Unless these channel realities are recognized and dealt with aggressively, it is impossible to deliver high QoS system performance even with a fully synchronous MAC layer. Indoor multipath and the frequency-selective fading that results is so complicated that a probabilistic-based system design is crucial. Simple buffering techniques are simply inadequate to deal with the difficulties presented by the indoor wireless channel. The true overall system QoS depends upon (i) the reliability with which the next time-slot is available without collision for the link in question and (ii) the PER for that time-slot.

MAC/PHY – Co-Design of these two key layers for optimal performance

The design of an OFDM radio that supports the 54 Mbps mode (64-QAM Rate= ¾) with negligible performance loss due to phase noise and other RF circuit imperfections is far more difficult than designing a more traditional single-carrier system using the same modulation type and code for the Gaussian channel. The primary reasons for this include:

- Since the OFDM system uses N-subcarriers, the symbol rate is correspondingly N-times smaller than the single-carrier system thereby making local oscillator phase noise requirements correspondingly more difficult;
- OFDM waveforms exhibit a peak-to-average power ratio (PAPR) that is considerably higher than single-carrier waveforms;
- In the context of direct-conversion architectures, I and Q phase and amplitude balance are <u>frequency-dependent</u> quantities that must be closely matched across the entire modulation bandwidth on a <u>bin-by-bin</u> basis for OFDM systems whereas single-carrier systems enjoy an averaging-benefit across the entire modulation bandwidth;

As a consequence of these and other factors, the (coded) residual bit-error rate that is deliverable by a radio transceiver system is not necessarily that low for the higher signaling rates (e.g., 36 Mbps, 48 Mbps, 54 Mbps), and it may of course get considerably worse when frequency-selective multipath is introduced into the picture. Everything possible must be done to ease the radio design requirements because the minimum requirements are already quite challenging and manufacturing device yields will drive the overall product cost if excessive requirements are levied.

One of the more significant factors that drive the required performance level (i.e., residual BER) of an OFDM radio is the data packet length used in the system. Assuming statistically independent bit errors and a 32-bit CRC field for each packet, the normalized payload throughput versus data packet length and bit error rate including FEC (CBER) is as shown in Figure 2. In 64-QAM OFDM operation, a CBER of 10^{-4} to 10^{-3} is desirable based upon radio-related design issues whereas these CBER levels clearly require packet lengths that are substantially less than 1500 bytes in order to avoid a substantial throughput efficiency penalty. As shown in Figure 3 for 64-QAM rate ¾ operation in an otherwise perfect 802.11a OFDM system, phase noise performance is very crucial in achieving CBER levels below roughly 10^{-3}. As evidenced by this figure, overall system phase noise (i.e., transmitter + receiver) must clearly be less than 2 degrees rms if the needed CBER performance level is to be approached based upon the allowable system PER.

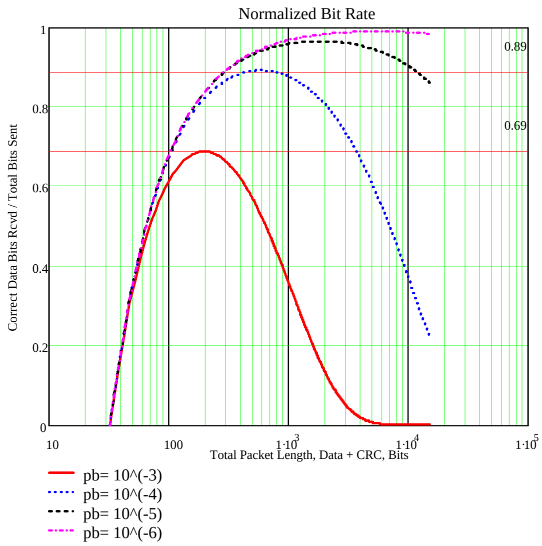

Figure 2: Throughput Efficiency Tradeoffs Involving BER and Data Packet Length

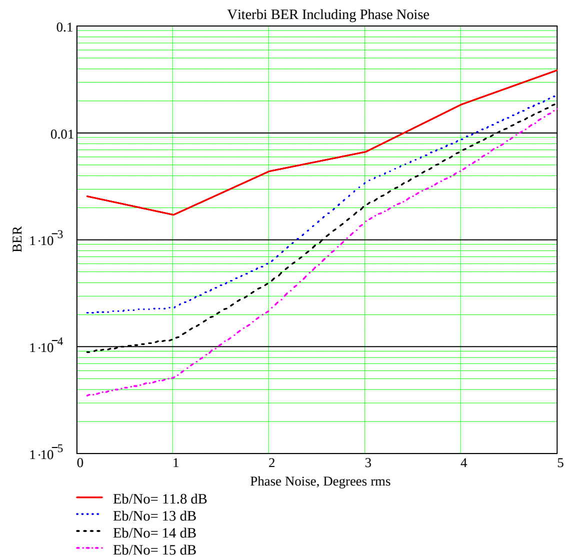

Figure 3: Air5 Security Diagram

To first-order, the system penalty for a non-zero PER is an increased throughput requirement level of $1/(1-PER)$ in order to deal with the packet re-transmissions involved, but in the QoS sense, the penalty is much more severe. If the PER is 15% for example, and only 3 packet transmission attempts are allowed in order to bound the QoS-related time-jitter, the probability that the packet is never delivered without error is still 0.34% which is generally not acceptable for video and audio. In a synchronous MAC system, more transmission attempts can generally be used because the time-jitter can be acceptably bounded whereas in asynchronous systems like 802.11x, the next time-slot availability question (without collisions) makes the argument for allowing more attempts much more dubious for high-QoS streams.

Many other factors favor synchronous MAC operation over asynchronous when it comes to radio design and digital signal processing as well. Since all network users are dynamically assigned time-slots, power-levels, etc., every receiver in the wireless network knows a priori the power-level, modulation-type, etc. This reduces linearity issues and estimation theory issues substantially in the receiver compared to an asynchronous collision-based system where no one knows what terminal will grab the next time-slot a priori. A synchronous MAC structure that allows periodic assessment of the RF channel (i.e., channel estimation) is also invaluable for dealing with frequency-selective multipath which is very major issue for indoor environments.

RF/PHY – You can't ignore multipath, you must embrace it.

The IEEE802.11x standards do not include any requirements nor provisions for multipath aside from the OFDM-nature of the physical layer waveform used. In contrast, anyone who has worked with high-throughput systems, whether OFDM or single-carrier, knows that frequency-selective multipath is a very serious issue for the indoor wireless environment.

OFDM alone is insufficient to deal with the indoor wireless channel when it comes to high-QoS applications like audio and video. Some of the reasons for this statement were mentioned under the System Design section of this paper. Realistically, it is impractical to deliver a recovered SNR through a low-cost transmitter-receiver pair in excess of approximately 30 dB due to phase noise and linearity issues. Therefore, if the system requires a minimum SNR of perhaps 24 dB at sensitivity for 54 Mbps operation, it is impossible to operate the link with video-adequate performance if there are any frequency-selective fading nulls greater than about 6 dB.

A large amount of recent research has been focused on space-time coding (STC) which primarily seeks to increase the throughput capacity of systems in a bits/Hz bandwidth perspective. Although this perspective is valuable in many venues, in the audio/video consumer realm where large data buffers translate into cost and time latency, it is difficult at best to use the additional throughput offered by such STC systems if the throughput rate is unreliable over time as is generally the case for an indoor channel. Just as STC systems exploit the multiple-input-multiple-output (MIMO) increased channel dimensionality that multipath channels exhibit to increase throughput capacity however, a measure of this same MIMO channel dimensionality can be similarly exploited to deliver dramatically more reliable link communications using OFDM signaling that is almost identical to IEEE802.11a. This probabilistic-based system design methodology which focuses on link reliability rather than on just average or peak throughput is extremely well-suited to the audio/video delivery needs of most applications.

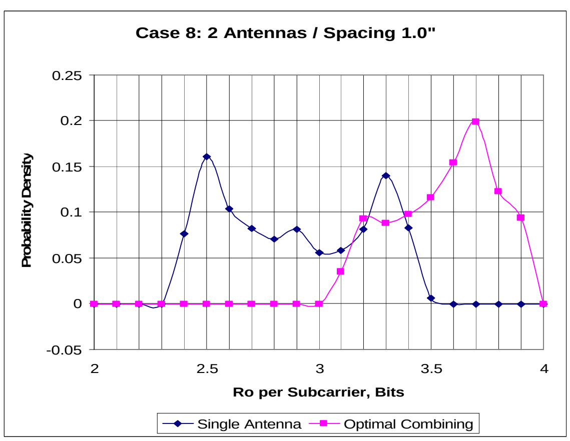

Figure 4: Two-Antenna Case

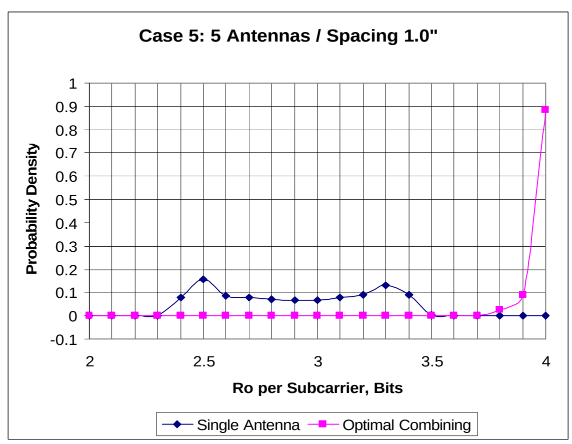

Figure 5: Five-Antenna Case

Computer simulations have confirmed using information-theory based metrics that under three-ray multipath conditions where a linear array of receive antennas is used assuming random direction-of-arrivals of the multipath rays that a minimum of 3 independent looks at the incident signal wavefront (i.e., 3 antennas) is required in order to deliver acceptable link throughput reliability for 16-QAM rate ¾ . The limited k=7 Viterbi FEC used in IEEE802.11a along with other practical implementation factors force the number of receive antennas to 4. When other real-world factors like outage probability and finite SNR are included, a larger number of antennas is still advantageous. The computer analysis was done based upon cutoff rate per subcarrier and the results for the (a) two-antenna and (b) five-antenna cases are shown in **Figure 4** and **Figure 5** respectively. As shown there, the single-antenna system delivers a wide span of throughput levels in 16-QAM ¾ mode (R_o= 2.3 to 3.5 bits/subcarrier/channel-use) because it has no means to deal with the frequency-selective fading introduced by the 3-ray multipath model. The use of two antennas improve the worst-case R_o to approximately 3 bits/subcarrier whereas the results dictated that five antennas be used in order to deliver the outstanding R_o performance shown in **Figure 5**. Additional details are provided in [3].

REAL WORLD TEST RESULTS

The amount of frequency-selective multipath that potentially interferes with reliable high-throughput high-QoS links is very severe making exhaustive field testing mandatory before a system design concept is truly validated. Multimedia-centric wireless links require the link reliability to be several orders of magnitude better than that commonly acceptable for data-centric wireless networks like 802.11x and this puts additional emphasis on the need for substantial field testing.

The spatial wavefront processing[1] advocated herein avoids the multipath-related losses that are common-place with indoor environments. Signal absorption losses are very environmental-dependent and Magis has therefore tested these concepts in a variety of homes, offices and business locations. Two home-testing trials are reported on separately in [4-6]. As reported in [4], full US-HDTV video[2] was delivered throughout a two-story 2,400 square-foot home as well as a second two-story 3,500 square-foot home. For data-networking applications, the throughput was consistently 40 Mbps or more in all but the most challenging areas of the homes. All of these measurements were made with the first-generation Air5 chipset whereas the second generation chipset exhibits substantially improved performance and is now sampling. Japanese home testing results in Figure 6 and **Figure 7** also provided excellent full-home coverage.

Figure 6: Japanese Home in Osaka Used for Trial ($240\ m^2$)

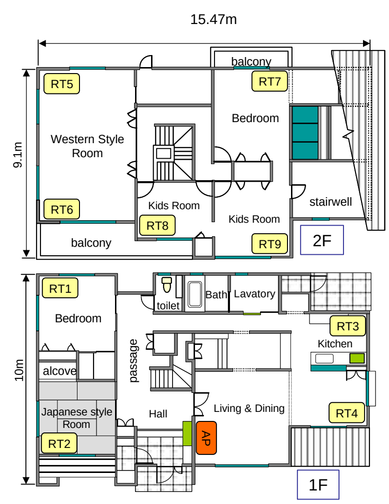

Figure 7: Full Japanese HDTV Coverage Delivered for the Entire Home

---

1 Referred to as CSWR within Magis

2 Video performance criteria was zero packet errors over a 2 minute measurement interval.

SPECTRUM ETIQUETTE

Spectrum etiquette is an important subject in the context of supporting high-QoS A/V links. While contention-based competition for time-slots per existing 802.11x standards has been adequate for data-centric applications, it only takes one non-QoS user to deny everyone on a given channel any chance for high-QoS performance. Unfortunately, 802.11x systems generally transmit at full-power regardless of the distance between network nodes thereby preventing re-use of the specific RF channel for very large distances except for possibly other non-QoS users. Everyone should care about QoS regardless of whether they are trying to network with 802.11x systems or with other more QoS-centric systems if they are trying to deliver A/V content in the home.

The FCC is moving rapidly to expand the useable 5 GHz spectrum in the US by 255 MHz. Usage of this spectrum carries with it certain caveats regarding other primary users like weather radar however, forcing yet another flavor of spectrum etiquette [7]. More stringent transmit spectrum mask requirements for all 5 GHz devices are advocated in [7] as well as other recommendations that will enhance everyone's opportunities in the 5 GHz U-NII bands.

CONCLUSION

True high-QoS support for emerging audio-video wireless networking applications requires a fundamentally different systems approach than traditional data-centric networking technologies because link reliability is equally as important as throughput rate. High-QoS wireless networking will enable a wide range of applications and products in the home and office. Co-design of the MAC, PHY and RF is required in order to achieve the end performance and cost objectives.

References[3]

1. ISO/IEC 13818-1: Information Technology – Generic coding of moving pictures and associated audio information: Systems, 1996
2. R. Wilson, J.A. Crawford, "Propagation Losses Through Common Building Materials", August 2002
3. J.A. Crawford, "Making OFDM Work for High-Performance Wireless Network Applications", April 2003
4. ________, "Air5 In-Home Field Testing Results", March 2003
5. ________, "3-D Tour of In-Home Testing: First Floor", March 2003
6. ________, "3-D Tour of In-Home Testing: Second Floor", March 2003
7. ________, "Comments Regarding FCC 03-110", 11 June 2003

[3] References [2-6] are available at www.magisnetworks.com

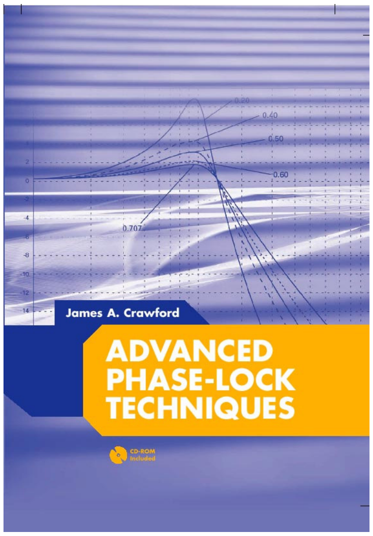

Advanced Phase-Lock Techniques

James A. Crawford

2008

Artech House

510 pages, 480 figures, 1200 equations
CD-ROM with all MATLAB scripts

ISBN-13: 978-1-59693-140-4
ISBN-10: 1-59693-140-X

Chapter	Brief Description	*Pages*
1	*Phase-Locked Systems—A High-Level Perspective* An expansive, multi-disciplined view of the PLL, its history, and its wide application.	26
2	*Design Notes* A compilation of design notes and formulas that are developed in details separately in the text. Includes an exhaustive list of closed-form results for the classic type-2 PLL, many of which have not been published before.	44
3	*Fundamental Limits* A detailed discussion of the many fundamental limits that PLL designers may have to be attentive to or else never achieve their lofty performance objectives, e.g., Paley-Wiener Criterion, Poisson Sum, Time-Bandwidth Product.	38
4	*Noise in PLL-Based Systems* An extensive look at noise, its sources, and its modeling in PLL systems. Includes special attention to $1/f$ noise, and the creation of custom noise sources that exhibit specific power spectral densities.	66
5	*System Performance* A detailed look at phase noise and clock-jitter, and their effects on system performance. Attention given to transmitters, receivers, and specific signaling waveforms like OFDM, M-QAM, M-PSK. Relationships between EVM and image suppression are presented for the first time. The effect of phase noise on channel capacity and channel cutoff rate are also developed.	48
6	*Fundamental Concepts for Continuous-Time Systems* A thorough examination of the classical continuous-time PLL up through 4^{th}-order. The powerful Haggai constant phase-margin architecture is presented along with the type-3 PLL. Pseudo-continuous PLL systems (the most common PLL type in use today) are examined rigorously. Transient response calculation methods, 9 in total, are discussed in detail.	71
7	*Fundamental Concepts for Sampled-Data Control Systems* A thorough discussion of sampling effects in continuous-time systems is developed in terms of the z-transform, and closed-form results given through 4^{th}-order.	32
8	*Fractional-N Frequency Synthesizers* A historic look at the fractional-N frequency synthesis method based on the U.S. patent record is first presented, followed by a thorough treatment of the concept based on Δ-Σ methods.	54
9	*Oscillators* An exhaustive look at oscillator fundamentals, configurations, and their use in PLL systems.	62
10	*Clock and Data Recovery* Bit synchronization and clock recovery are developed in rigorous terms and compared to the theoretical performance attainable as dictated by the Cramer-Rao bound.	52