Implementation of an Experimental 384 kb/s Radio Link for High-Speed Internet Access

Bruce McNair, Leonard J. Cimini, Jr., and Nelson Sollenberger

AT&T Labs - Research,
100 Schulz Drive, Red Bank, New Jersey 07701-7033

e-mail: bmcnair, ljc, nelson@research.att.com

ABSTRACT

Interest in wireless high-speed access to wide area data networks has grown considerably in recent years, fuelled by the simultaneous development of the high-speed wired data networks, wide-spread availability of powerful mobile computing devices, rapidly dropping prices for increasingly sophisticated electronic devices, plus the global deployment of wireless voice systems. For many users, high-mobility wireless data access at near-LAN rates is likely to be an attractive option. Unfortunately, the harsh radio environment creates significant challenges when trying to provide high-speed access, especially if high-mobility and wide-area availability are expected. This paper describes an experimental implementation of a 1900 MHz OFDM radio link designed to provide 384 kb/s data link in representative mobile cellular-like channels.

1. INTRODUCTION

The interest in wireless high-speed access to wide area data networks has grown considerably in recent years. This growth is fuelled by the simultaneous development of the high-speed wired data networks (e.g., the Internet), significant improvement in data network access speeds (e.g., cable modems and DSL), wide-spread availability of powerful mobile computing devices (e.g., PDAs and notebook computers), rapidly dropping prices for increasingly sophisticated electronic devices, plus the global deployment of wireless voice systems. There is no reason to expect this momentum to decrease in the foreseeable future. For many users who are beginning to experience what high-speed global data access or untethered voice access can mean for them, the jump to wireless data access at near-LAN rates is attractive. Unfortunately, the harsh radio environment, with fading, dispersion, noise, limited spectrum, power restrictions, and a host of other constraints creates challenges when trying to provide this high-speed access, especially when the high-mobility seen in today's cellular networks is expected.

Investigations into wireless communications systems for Third Generation wireless systems (3G) and beyond have dealt with a wide variety of communications techniques. Our interests in this area have included OFDM[1][2] for reasons that will be set forth below. While idealized computer simulations of such systems can provide valuable insight into their operation, many impairments and system behaviors can only be studied effectively with a realistic, real-time implementation. Simulations can be criticized for their optimistic assumptions, e.g., perfect knowledge of channel response, frequency offset or timing, to name a few. As simulations are made more and more realistic to counter these criticisms, the amount of processing on a general purpose workstation or compute server needed to model and respond to practical impairments leads to excessive simulation times. Finally, simulations can best be used for performance measurements under a limited set of channel conditions or user traffic models. To demonstrate the performance a user could expect under actual traffic conditions and on realistic channels requires a real-time implementation. This paper describes one such implementation of a system proposed for use as the downlink for high-speed (384 kb/s and faster) Internet access.

We have chosen to restrict our study to the downlink for two important reasons. First, across a large aggregation of users, as well as for most smaller sets of users, asymmetry in traffic levels is prevalent, with most user terminals receiving much more traffic than they transmit. High-volume information providers are usually best served by fixed, wired connections while mobile users are most likely to be information consumers. Second, existing and soon-to-be-fielded wireless data services provide more symmetric capacities. We expect that a higher speed downlink could be an overlay network, with the lower capacity symmetrical networks used for uplink, control, signalling, and a downlink for smaller information units.

The remainder of this paper is organized as follows: First, a general overview of OFDM is presented. We then present the specific OFDM parameters chosen for this implementation and describe the system architecture, both from a hardware and software perspective. We present the specific modulation and coding parameters that were used in the prototype and describe the experimental setup. We conclude by describing the future directions for this research. Experimental performance results are presented in a separate paper[3].

2. OFDM BACKGROUND

Delay spread (time dispersion of a brief transmitted pulse) is a major impediment to high-speed data transmission in the outdoor, high-mobility wireless environment. Distant signal reflections from terrain, buildings, or moving vehicles may create several microseconds of delay spread. The high-mobility (i.e., high vehicle speeds) environment also creates a rapidly changing channel. While adaptive channel equalization techniques can be used to combat some of this effect they are limited in their ability to deal with large amounts of rapidly changing delay spread, requiring longer, more complex equalizers. Conversely, using low symbol rates to mitigate channel dispersion avoids the need for complex equalizers, but requires very dense constellations in single carrier systems, increasing susceptibility to noise, interference, and fading. CDMA RAKE receivers mitigate the effects of delay spread by attempting to estimate and combine the channel rays individually. They do so at the cost of the high start-up bandwidth required of a CDMA system and rely on specialized signal processing hardware, among other constraints.

In an OFDM system, a single high-speed bit stream is separated into several low speed bit streams, with each of the slower bit streams used to modulate one of a set of carriers. The carriers are generated within a single bandwidth, with their inter-tone spacing chosen to insure orthogonality of the resulting combined waveforms. By designing the transmission in this way, the benefit of low symbol rates is attained without the penalty of dense constellations – each individual carrier has a sparse constellation, often QPSK. The union of a large number of these individually sparse constellations determines the total capacity of the OFDM system.

One could build an OFDM modem by generating a multitude of sinusoidal carriers and modulating them individually. However, a modern OFDM transmitter is typically implemented by generating a series of complex numbers, each complex number representing the phase and amplitude of the individual tones. Using the Inverse Discrete Fourier Transform (IDFT), or to reduce complexity, the Inverse Fast Fourier Transform (IFFT), this series of tones is converted into a time domain waveform. For practical reasons, in order to make the analog reconstruction and antialiasing filters realizable, the length of the IFFT is usually somewhat larger than the number of tones. Unused tones at the edge of the band are set to zero to eliminate energy outside the desired bandwidth.

Similarly, the OFDM receiver uses a forward FFT to transform the received waveform back to a set of orthogonal tones, ignoring the unused tones. The complex numbers that represent the received tones' amplitudes and phases directly correspond to the modulated carriers, allowing simple demodulation of the individual carriers. In many cases, OFDM can achieve equivalent performance to other modulation techniques at equal or lower complexity[6], due in large part to the signal processing efficiency and regular structure of the FFT. The regular structure of the FFT also lends itself well to signal processing on general purpose DSPs and other devices. Finally, unlike other forms of modulation, OFDM offers the greatest flexibility in bandwidth assignment. If more or less spectrum is available for an OFDM-based service, the selection of transmitted tones can be readily adjusted to best match the bandwidth available. Likewise, if certain discrete frequencies in a band are known to have high co-channel interference or must be avoided for other reasons, the OFDM system could be readily configured to avoid those frequencies while using the rest of the band.

In addition to the unused tones at the IFFT input, extra samples are prepended and appended to the transformed waveform. These "cyclic prefix" and "cyclic suffix" samples make the transmitted signal more robust against time dispersion and timing offset. As long as the received signal, plus its various delayed copies, are sampled within the FFT plus cyclic extension interval, the constant amplitude of the received spectrum will be preserved. By the time-shifting property of the FFT, delayed copies of the transmitted waveform appear as signals in the frequency domain with frequency dependent phase rotation. . We took advantage of this latter fact in designing an extremely robust timing recovery algorithm, previously described[4]. As we show in the related paper on system performance, negligible performance degradation is seen for any level of delay spread within the cyclic extension interval.

Figure 1 shows a simplified block diagram of an OFDM transmitter. Note that the output is the complex baseband signal which will be translated to RF by an IQ modulator. For this prototype, all processing, both in the time domain and the frequency domain, was performed on the complex representation of the signals, generally requiring two parallel paths for the real and imaginary parts of the signals.

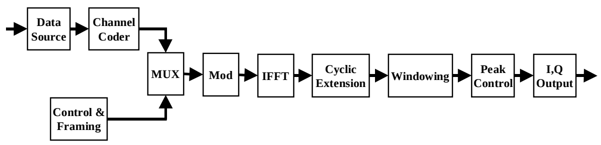

Figure 1 - Block diagram of OFDM transmitter

For the OFDM modem described in this paper, modulated tones occupied about one-third of the total number of tones. Because of the structure of the FFT, using a total processed bandwidth that is a power of two times the tone spacing provides the most efficient signal-processing algorithm. The number of active (i.e., modulated with user data and signalling) tones is chosen to provide the desired bandwidth and data rate for the system. Therefore, it would generally be desirable to pick the next higher power of 2 for the FFT size. In our case, with 189 active tones, a 256 point FFT would have lessened the real-time signal processing burden by about 55% compared to the 512 point FFT that was used. However, the smaller FFT would have resulted in a baseband signal with a stop band that began about 1.5 times the highest desired frequency. Such a filter is realizable. However, when considering the filter complexity and the inevitable group delay distortion, choosing the larger FFT was the simplest choice. This allowed use of a filter that has a stopband frequency more than 4 times the highest desired frequency. Obviously, for a real product, the cost of signal processing versus the cost of designing, manufacturing, and testing the antialiasing and reconstruction filters would have to be traded against each other. With digital device processing speeds increasing exponentially (with concomitant reductions is per instruction processing costs) while analog component and manufacturing costs remain level or are increasing, choosing the larger FFT is likely to remain the preferred option.

Two other segments are added to the IFFT waveform: the guard interval and the windowed section. The guard interval is a period during which the transmitted samples are all set to zero. At the expense of a slightly reduced transmission efficiency, the presence of this guard interval helps to insure that the samples received during a given OFDM block are not contaminated with delayed samples from a previous block or from a distant base station. Since the OFDM system under study is used in packet mode, it is expected that the transmitter will be turned on and off during each OFDM burst. To prevent the radiated signal from "splattering" outside of its assigned channel due to a rapid turn-on or turn-off of tone energy, "windowing" of the OFDM samples is used. By shaping the transmitter ramp up and ramp down with a raised cosine pulse shape, the system bandwidth is confined to little more than the bandwidth occupied by the set of OFDM tones themselves. Windowing also reduces the system's sensitivity to frequency offset and Doppler by reducing the amplitude of the interference contribution from adjacent tones. Figure 2 illustrates these items and provides the particular parameters used for the prototype.

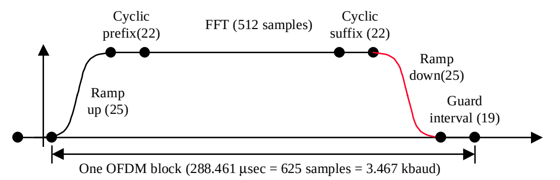

Figure 2 - OFDM block structure

The OFDM receiver structure mirrors the operation of the OFDM transmitter. Again, time domain and frequency domain signals are treated as complex. In the receiver, a second form of windowing is used, where samples in the cyclic extension are combined with the samples corresponding to the IFFT output samples at the transmitter. Receive windowing will improve rejection of adjacent channel interference and may aid performance against co-channel interference and other impairments. Figure 3 shows the means by which receive windowing is performed. Since, at the transmitter, the cyclic extension samples before (or after) the FFT interval are copies of the FFT samples at

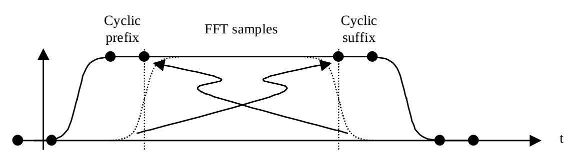

Figure 3 - Receive Windowing

the end (or beginning) of the FFT interval, assuming proper timing and frequency offset correction at the receiver, these received samples should be identical and can be directly combined. The receiver samples at the edge of the FFT interval are reduced in amplitude and added to the corresponding cyclic extension samples, similarly reduced in amplitude to yield the correct final amplitude. With a windowing function f() defined on the interval [-k,k], a set of received samples r_i, and an N-point FFT, the set of samples, w_i, after windowing is:

$$w_i = f(i) \cdot r_i + \{1 - f(-i)\} \cdot r_{N+i} \qquad (1)$$

Finally, the extra samples are removed from the received waveform and the resulting N time samples are transformed into the frequency domain. Here, assuming that there was no frequency or timing offset in the time domain waveform, a series of tones are processed, with each conveying one of the lower speed bit streams. If QPSK were used at the OFDM transmitter to modulate each tone, then each OFDM tone burst would convey two bits of information per tone in the phase of that tone. A simplified OFDM receiver is shown in Figure 4.

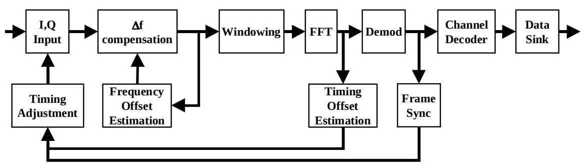

Figure 4 – Block diagram of OFDM receiver

While the relatively long symbol period of the OFDM waveform makes the system resistant to delay spread, incorrect timing phase or frequency offset in the received signal will quickly degrade system performance. Most of the extra samples in the time domain waveform insure that the amplitude of the received OFDM spectrum remains constant in the presence of delay spread and timing offset, due to the cyclic shifting property of the Fourier Transform. However, information is conveyed in the phase of the received tones, so the damaging effects of timing offset must be corrected. After estimating the amount of timing offset in the received signal using the method previously, the phase of the sample clock was continually adjusted to drive the offset to zero.

Frequency offset has a different effect on the received waveform. While the OFDM transmitter generated orthogonal tones by creating a specific linear harmonic relationship between all the tones, frequency offset shifts each tone by an equal amount. Thus, the frequency shift creates an affine relationship between tone frequencies, destroying the linear harmonic structure and creating inter-channel interference (ICI). Like inter-symbol interference in a single carrier system, frequency offset allows energy from adjacent tones to "bleed" into a desired tone, reducing system performance. Frequency offset was compensated for by performing a complex multiplication between the received samples and a complex vector rotating at the estimated offset frequency.

3. OFDM PARAMETERS

As an example of a practical system, the OFDM radio link described here is intended to provide peak end user data rates of at least 384 kb/s. As shown in Figure 2, this system transmits OFDM bursts with a duration of 288.462 μsec. The sample rate, the OFDM burst duration, and the 30 msec OFDM frame are derived from the GSM-based clock frequency. In addition, the choice of the burst duration is a tradeoff between the fading rate on the channel (shorter blocks), the expected delay spread (longer blocks), and the transmission efficiency (fewer redundant samples per FFT sample).

The 512 FFT samples used here represent the transform of 189 discrete complex tones. These 189 tones are spaced every 4.232 kHz (skipping the three tones nearest the center frequency to avoid a DC component) so the total bandwidth is 812.5 kHz. With QPSK modulation on each complex tone, this means that the raw channel rate without coding is 1.3104 Mb/s. With a rate-½ Reed-Solomon channel code plus a generous allowance for framing and control overhead, the peak end user capacity goal of 384 kb/s is easily attained.

4. SYSTEM ARCHITECTURE

Although OFDM modulation mitigates the effects of time dispersion of the transmitted waveform, there is another effect due to the dispersive channel that must be combated – frequency selective fading. Consider the simplest case where the dispersive fading channel can be represented as two independently fading rays with a constant time separation. If each ray is considered to be a Rayleigh fading signal, at times the rays will have the relative phases and amplitudes to cause nearly complete destructive interference at some frequencies. If the channel response is transformed into the frequency domain, it can be seen that there are frequencies where the received signal will be so attenuated that it may be useless for data transmission purposes. However, a second receiver, with a separate antenna, will receive a slightly different set of signals. The attenuation and center frequency of spectral nulls is quite sensitive to the amplitude and phase of the rays. Thus, even with moderate amounts of correlation between the antennas' signals, it is unlikely that a second antenna will receive *exactly* the same signal energies at the *same* frequencies for any appreciable period of time. For this reason, OFDM is particularly well suited to being combined with antenna diversity. If the delay spread does not extend beyond the cyclic extension, the receiver is robust against the time dispersion effects. As long as the echoes that are received are not all too close together in time, the mostly independent frequency nulls will allow the receiver to process OFDM tone energy at all frequencies. Thus, it was recognized early in this investigation that antenna diversity, which is generally valuable for all wireless communications techniques, should prove to be quite beneficial in conjunction with the inherent robustness of OFDM.

One other feature of the channel frequency response should be noted. At any given frequency, severe fades may be present. Due to the mechanism of the fading, that is, destructive cancellation of the signal due to the phase and amplitude relationships of the fading rays, it is easy to see that fades will tend to influence signals over a relatively narrow range of frequencies, depending on the relative delays between paths. However, most of the frequencies are passed with little or no degradation. Because energy is conserved, we can see that there are many frequencies where the instantaneous energy is actually increased, giving a higher SNR and better error rate. We can take advantage of this fact if we code across the band of frequencies, averaging the effects of frequency selective fading. Frequency-domain coding provides an analog of time-domain interleaving. What is most important is that on slow fading channels this method does not create the inherent, often unacceptable, time delays that an interleaver would. For bursty traffic in a packet transmission system, interleaving delays could significantly degrade system performance. Only for "flat fading" or for channels that have very closely spaced (in time) dominant rays will the frequency diversity coding advantage be lost.

As part of our investigations, several methods of antenna combining, channel estimation, and interference suppression have been studied for OFDM systems[7][8]. For the purposes of this initial real-time prototype, a simple demodulation/combining strategy was used. Data symbols on each individual carrier were differentially encoded from one block to the next in time. Thus, the receiver uses the set of tones transmitted 288 μsec earlier to demodulate the phase of the current set of tones. While this incurs a theoretical 3 dB penalty in the decoding of the individual tones compared to a coherent receiver in AWGN, the greatly simplified receiver structure made the prototype much more practical to implement. The combining technique used was equivalent to maximal ratio combining – initially, in steady state conditions, the gain of two receivers was matched at the demodulators inputs by measuring and balancing the measured aggregate tone powers. Since each channel had an equivalent differential phase reference from its own delayed samples, the two demodulator outputs were simply added. Adding the two channels' constellations in this way provided an inherent receiver branch weighting based on the reliability of each signal's validity. As demonstrated in our paper on system performance, an OFDM system performs best in the delay spread environments that degrade single carrier systems. Two receiver branches with straightforward combining provide substantial performance benefit over a single receiver. Thus, delay spread, combined with simple antenna combining creates very effective frequency diversity that, in combination with channel coding, improves the overall performance.

As stated, the modulation was performed differentially in time. While maximal ratio combining significantly reduces the impairment of frequency selective fading, it can be expected that there still would be degradation in SNR at certain frequencies. Since we envision that the OFDM system will need to operate at very low average signal-to-noise ratios, we decided to use Reed-Solomon coding across the tones in the OFDM block. The Reed-Solomon code that was available to us was designed for GF(64) with 6 bits per code symbol and 63 channel symbols per code word. With two bits per QPSK-modulated tone, this meant that three tones corresponded to one Reed-Solomon symbol and the 63 Reed-Solomon symbols in a code word corresponded to 189 tones in an OFDM block.

As we have shown, on a dispersive channel, it is likely that some tones will be received with a lower signal power and, thus, probably a poorer SNR than other tones. An (n,k) Reed-Solomon code[9] can correctly decode a block with ρ erasures and ν errors, as long as $\rho + 2\upsilon \le n - k$. We felt it would be advantageous to use some of the power of the code to treat received symbols as erasures where there was good reason to suspect the validity of the symbols. For this prototype, tones were collected into groups of three that corresponded to the Reed-Solomon symbols and the total power of these groups was recorded as a measure of symbol reliability. With all 63 Reed-Solomon symbol powers known, the smallest *m* were arbitrarily marked as erasures. While the optimum value for *m* varies with channel conditions, a value of *n*/4 (i.e., half of the "power"of the code used for erasures and half used for error correction) was found experimentally to be a reasonable choice across all channel conditions.

5. DSP IMPLEMENTATION

We have built high-speed signal processing systems based on the Texas Instruments TMS320C40 floating point DSP for several prior research programs; in particular, systems based on general purpose DSP boards from Pentek. Naturally, the prior experience with this line of products made it a logical choice when we started the OFDM program. With the goal of building a two-branch 800 kHz OFDM receiver on the 'C40-based signal processing platform, it became obvious that multiple DSPs would be required to implement all the receiver functions. Figure 5 illustrates the final assignment of functions to individual DSPs as well as showing the logical interface to other hardware.

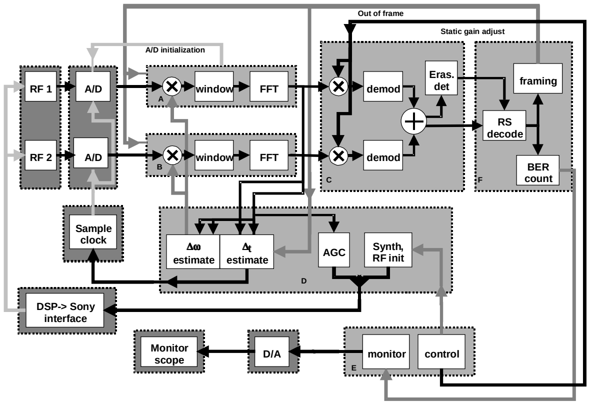

Figure 5 - OFDM Modem Signal Processing Block Diagram

Individual DSPs were dedicated to each of the following high-level functions:

- **FFT front end processing**, including acquiring samples from the A/D converter at a 2.166 MHz sample rate, baseband rotation to compensate for frequency offset, receive windowing, and calculating 512 point complex FFTs. This functional block also included calculating estimates of the frequency offset of the received signal. One DSP was used per receiver branch for this function.
- **Demodulator back end processing**, including static gain balance between receivers, differential demodulation, combining, and erasure detection
- **Decoder processing**, including the Reed-Solomon decoder, framing detection, and maintaining frame error rate statistics
- **Signal control processing**, including sending gain and frequency control words to the external RF unit, measuring receive signal power levels to perform automatic gain control, estimating timing offset and sending timing and frequency offset corrections to functional blocks that could act on these instructions. This included sending clock add/delete instructions to a clock board designed for the OFDM project.
- **Control monitor functions,** including real-time user control of receiver parameters and allowed real-time display of various signals throughout the receiver.

Essentially all of the interprocessor communications for the receiver occurred via shared global memory access. This allowed a looser coupling between individual processes than might be required had the ‘C40s high-speed communications ports or other techniques been used. With shared memory access, the sending process wrote its data to its global memory as needed. When a buffer was complete, a semaphore was sent, again, via the shared global memory. The receiving process could perform its operation asynchronously with the sending process until it needed another buffer. At this time, it began polling the semaphore, waiting for a signal that the sender had finished. By insuring that no downstream process ran slower than the front end FFT processes, it was possible to guarantee that no blocks were lost in the chain.

With the Pentek 4270s that were used, four ‘C40s were available on each board. The receiver functions that required the greatest intercommunications, namely the FFTs, the demodulation, and the signal control functions, were placed on the same 4270, minimizing the off-board communications bandwidth needed. Since the decoder only received 63 6-bit Reed-Solomon symbols per OFDM block plus no more than half that much information in the form of erasure locations, it was reasonable to transfer detected symbols over the Pentek MIX bus[10]. Likewise, although the control/monitor function accessed and controlled real-time processes, the user inputs and displays could be operated at significantly lower update rates, allowing both of these functions to be assigned to the second 4270 board.

Of course, with any multi-processor system, synchronization of the state of each processor is always a challenge. The complexity of this design was increased since there were two sets of A/D channels (I and Q for each receiver branch). In addition, at a 2.166 MHz sample rate, it was not feasible to require that one DSP handle both branches. Thus, two DSPs, individually acquiring samples, had to synchronize within a less than a microsecond, while the only communications mechanism between the two processors had more than 200 nsec of one-way latency. That, plus the fact that the ‘C40 has no hardware semaphore mechanism, meant that all of the signalling had to be carried out in software. One final constraint was the need to maintain a reasonable level of signal processing in the front end DSP while acquiring OFDM block samples at a high rate. While operating an interrupt service routine (ISR) to oversee the ‘C40’s DMA (direct memory access) from the comm port to a memory buffer had minimal impact on the performance of the background FFT processing, this meant that there were now four individual processes running on the two DSPs. Figure 6 illustrates the interprocessor signalling functions that were used to insure consistency across all six receiver DSPs. Each of the larger blocks corresponds to a different DSP. It should be noted that the two blocks in the upper left, the FFT functions, are actually executing identical code. Based on whether the FFT process found itself on the first or second DSP, it ran a slightly different subset of the code. Using identical software for shared functions made it easier to maintain the critical functions.

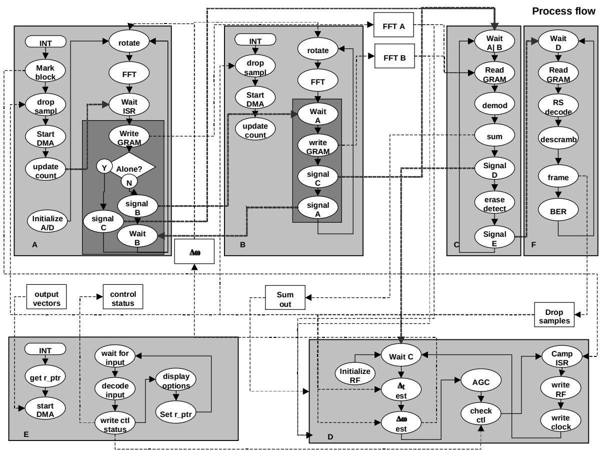

Figure 6 - OFDM Receiver Process Synchronization

6. OTHER IMPLEMENTATION DETAILS

The bulk of the effort for this prototype was writing and debugging the real-time C language software. Other than a small fraction of the code that consumed a substantial fraction of the total real-time (e.g., the FFT and a number field transform for the channel coder), the use of a high-level language like C made implementation much simpler. Since the hardware signal processing platform was mainly assembled from commercial off-the-shelf components, hardware integration, although not an insignificant task, was much simpler. There were, however, four hardware components of the system that could not be obtained commercially. These components are described below.

First, since we wanted to be able to realize actual timing recovery instead of artificial, idealized, "timing under the table," we needed to design and build a clock generation board to drive the sample rates of the A/D and D/A converters. Clock boards were designed that generated the 2.166 MHz sample rate used by the receiver and transmitter, as well as OFDM block and frame clocks that can be used for synchronizing scope displays of the waveforms. The 78 MHz master clock for the clock board is a multiple of the 13 MHz master clock frequency used in GSM systems. For receiver timing frequency and phase adjustment, the ability to advance or retard clock phase in 25 nsec intervals was provided, under the control of the DSP's high speed comm port. An Altera field programmable gate array (FPGA) was used to provide the comm port interface as well as perform the necessary clock division and control logic.

Next, we needed a high-speed D/A converter. Since most DSP board vendors have developed their business around the needs of the military and intelligence community (who mostly listen, and rarely talk), there are few D/A converters with the resolution, conversion rates, or interfaces we needed. An eight channel D/A converter was designed that used the same 'C40 comm port interface and could be plugged into the signal processing hardware platform's VME bus. Each comm port is theoretically designed to operate at a 5 MHz sampling rate. With a conservative approach to signal buffering, a 3 MHz sample rate was reliably attained. Each comm port delivers 32 bit samples, consisting of 16-bit I samples and 16-bit Q samples. Thus, the D/A could provide two transmitter branches and a constellation display output with one spare IQ output available.

At the transmitter, the D/A samples must be low-pass filtered to reconstruct the baseband analog signals. Likewise, the receiver requires low-pass filtering to prevent high frequency noise and interference from aliasing down to the receiver's passband. As noted above, there is interaction between the sample rate, the FFT size and the analog filter requirements. Since the 800 kHz desired bandwidth of the transmitter was given, a filter was required that had little attenuation or group delay at or below 400 kHz, the highest frequency in the complex baseband signal. Likewise, the first alias frequency can be found, knowing the sample rate and the complex signal bandwidth. For the parameters chosen, the first alias frequency is at 2166 – 400 kHz = 1766 kHz. A commercially available filter module was chosen with an 8-pole Butterworth response and a 3 dB frequency of 550 kHz. This design insures that the attenuation at the end of the passband (here, 73% of the corner frequency) is less than .01 dB with less than 200 nsec of group delay variation across the passband. At the first alias frequency (here, 3.2 times the corner frequency), the filter response provides at least 82 dB of attenuation. Since the low-pass Butterworth amplitude response is monotonically decreasing, this insures all other alias signals are attenuated at least 82 dB. In fact, at most alias frequencies, the attenuation is considerably higher.

Finally, to allow operation through a RF fading simulator or over a real outdoor channel, RF hardware was needed that would convert between baseband I&Q signals and a 1900 MHz RF signal. As a result of a collaboration with Sony on their OFDM-based "Band Division Multiple Access" (BDMA) project[11], they supplied us with mobile and base station transceivers. These transceivers were designed to operate across the 1900 MHz PCS bands with firmware controllable operating frequencies, transmitter and receiver gains, and RF bandwidths of 200 kHz or 1.6 MHz. Our 800 kHz OFDM signal readily fit in the wider bandwidth of their transceiver. Finally, Sony's transceivers provided options for one- and two-branch diversity for both the transmitter and receiver.

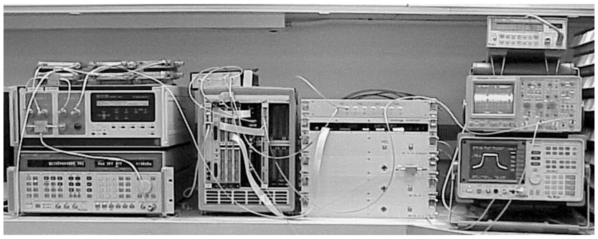

Figure 7 - Prototype base station with channel fading simulator

Figure 7 shows the base station prototype. From left to right are:

- The RF channel fading simulator, with an HP synthesizer serving as the fader's local oscillator and attenuators to set levels
- The DSP baseband signal processing hardware, including clock and D/A boards. The small box on top of the nest is a two section analog filter board.
- The Sony-provided RF transceiver, plus the top two modules we designed provide the baseband interface
- An oscilloscope displaying the OFDM waveform and a spectrum analyzer displaying the 800 kHz RF signal

7. CONCLUSIONS

This paper has described the architecture, design, and implementation of an OFDM modem intended for high-speed data access in wide-area, high-mobility applications. The experiments have demonstrated the practical considerations in OFDM system design and have yielded results that are being used to guide and validate decisions for the next generation, higher speed system.

As mentioned above, the parameters for this 384 kb/s OFDM modem were chosen to be comparable to those of the EDGE system as it was proposed for the 3G standards. Since the end user experience is highly dependent on the *peak* data rate that is available to a population of users we feel that it is necessary to consider higher capacity channels for future (e.g., beyond 3G) wireless data access networks. Our next phase of OFDM investigation focuses on 2-10 Mb/s in flexible 1-5 MHz channel assignments, building on the work presented here. To attain good performance at reasonable SNRs, of course, receiver antenna diversity and forward error correction coding will be required. In addition, we anticipate that channel estimation, interference cancellation, more sophisticated coding techniques, as well as adaptive modulation techniques will be required.

REFERENCES

[1] L.J. Cimini and N.R. Sollenberger, "OFDM with Diversity and Coding for Advanced Cellular Internet Service," *Proc, IEEE Globecom'97*, pp. 305-309, November 1997.

[2] J.Chuang, L.J.Cimini, G.Li, L.Lin, B.McNair, N.R.Sollenberger, M.Suzuki, H.Zhao, "High-speed wireless data access based on combining EDGE with wideband OFDM," *IEEE Comm. Magazine*, November 1999.

[3] B.McNair, L.J.Cimini, N.R.Sollenberger, "Performance of an Experimental 384 kb/s 1900 MHz OFDM Radio Link In a Wide-Area High-Mobility Environment," *Proc. VTC2000*, Boston, MA, September, 2000.

[4] B.McNair, L.J.Cimini, N.R.Sollenberger, "A Robust Timing and Frequency Offset Estimation Scheme for Orthogonal Frequency Division Multiplexing (OFDM) Systems," *Proc VTC99*, pp.690-694, Houston, TX, May, 1999.

[5] T.Kunihiro, T.Yamaura, M.Suzuki, "BDMA Testbed – Configuration and Performance Results," *Proc. VTC99*, pp.1886-1890, Houston, TX, May, 1999.

[6] H.Zou, H.J.Kim, S.Kim, B.Daneshrad, R.Wesel, W.Magione-Smith, "Equalized GMSK, Equalized QPSK, and OFDM, a Comparitive Study for High-Speed Indoor Data Communications," *Proc. VTC'99*, pp.1106-1110, Houston, TX, May, 1999.

[7] Y.(G.)Li, L.J.Cimini, and N.R.Sollenberger, "Robust channel estimation for OFDM systems with rapid dispersive fading channels," *IEEE Trans. Comm.*, vol. 46, pp. 902-915, July 1998.

[8] Y.(G.) Li and N.R.Sollenberger, "Adaptive antenna arrays for OFDM systems with co-channel interference," *IEEE Trans. Comm.*, vol. 47, pp. 217-229, February 1999.

[9] S.B.Wicker, V.K.Bhargava, *Reed-Solomon Codes and their Applications*, IEEE Press, NY, NY, 1994.

[10] *Operating Manual, Model 4270, Quad TMS320C40 Digital Signal Pocessor, VMEbus Board/MIX Module*, Pentek, Inc., Upper Saddle River, NJ, 1997.

[11] T.Kunihiro, T.Yamaura, M.Suzuki, "BDMA Testbed – Configuration and Performance Results," *Proc. VTC99*, pp.1886-1890, Houston, TX, May, 1999.