A 20 Mbits /s OFDM Demonstrator At 5 GHz: System Design, Implementation and Experimental Results

Robert Castle, Alan Jones
Home Communications Department
HP Laboratories Bristol

HPL-98-24
February, 1998

OFDM

The design and construction of a 20 Mbits/s OFDM demonstrator operating at 5.25 GHz in an indoor office/laboratory environment is described. The results from a field trial of this system are presented, from which the importance of coding in OFDM is illustrated and the degree of correctability required to achieve a 1% packet failure rate is derived.

A 20 Mbits/s OFDM Demonstrator at 5 GHz: System Design, Implementation and Experimental Results

R. Castle and A. Jones

Hewlett-Packard Laboratories, Bristol, UK

Abstract – The design and construction of a 20 Mbits/s OFDM demonstrator operating at 5.25 GHz in an indoor office/laboratory environment is described. The results from a field trial of this system are presented, from which the importance of coding in OFDM is illustrated and the degree of correctability required to achieve a 1% packet failure rate is derived.

I Introduction

Orthogonal Frequency Division Multiplexing (OFDM) has attracted considerable interest in recent years as a means of delivering high data rates with strong resistance to Inter Symbol Interference (ISI). It is currently being considered as a solution for Wireless ATM [1] and has previously been suggested for indoor applications where mobility and a high data rate are important, such as in [2]. It works by dividing a single high rate data stream into a set of parallel lower rate streams which are then transmitted together on a comb of orthogonal carriers. The inverse Fast Fourier Transform (FFT) is a convenient and efficient process for doing this, with the further benefit that its reciprocal process, the forward FFT, can be used in the receiver to recover the data.

In a frequency selective fading channel some of the carriers in the OFDM signal will be affected more than others, producing localised errors. Encoding the data prior to distributing it across the carriers provides each OFDM symbol with some redundancy. If sufficient redundancy is introduced then the errors from the few faded carriers can be eliminated from the symbol as a whole. The key issue is the degree of correctability required and hence how much redundancy must be added.

To investigate this, an OFDM test bed has been constructed which transmits a raw data rate of approximately 20 Mbits/s at a centre frequency of 5.25 GHz. The receiver captures bursts of signal corresponding to a packet-worth of data and then uses Digital Signal Processing (DSP) to implement symbol and frame synchronisation prior to demodulation, differential decoding and error detection. Statistics from this system have been collected during trials in an office/laboratory environment, looking particularly at the maximum number of errors occurring in an OFDM symbol during a packet. This gives the level of correctability required for that packet to have been received error-free. The distribution of these statistics therefore gives an indication of the packet failure rate which can be expected for a given coding correctability, or alternatively, the lower bound of the correctability required to achieve a certain packet failure rate.

The remainder of this paper is structured as follows: sections II and III describe the system design and implementation, section IV gives an outline of the experimental method used in the field trial, section V presents and discusses the results and the conclusions are in section VI.

II System Design

The demonstrator was originally conceived as being an experimental trial of a 20 Mbits/s OFDM system, where 20 Mbits/s is the uncoded throughput; clearly, with coding the user data rate would be less. On the basis of measurements and simulations [3], the system parameters were set out as follows:

- Measurements of the environment suggested that RMS delay spreads certainly greater than 50 ns and possibly as high as 100 ns were likely. As simulations have suggested that an irreducible error floor persists even with guard periods as large as 7 × the RMS delay spread, the guard period was set to be 1 μs.
- The guard period obviously has an effect on the data throughput of the system, since it is 'dead time'. 80% of the maximum was considered to be the worst degradation that would be allowed, making the OFDM symbol period 4 μs.
- If the desired 20 Mbits/s is achieved with a throughput of 80%, then the actual rate at which the system runs must be 25 Mbits/s. 25 Mbits/s in 4 μs means that 100 bits per symbol are transmitted, which for differential QPSK on each carrier requires 50 carriers.
- 50 carriers fits comfortably into a 64 point FFT with room to spare for filtering, plus any extra carriers which may be added in future as pilots for synchronisation *etc.* 64 samples in 4 μs makes the sample frequency 16 MHz, and the guard period therefore uses 16 samples.
- Having by nature a rectangular symbol shape, OFDM tends to have rather high out of band emissions unless the number of carriers is very large. To restrict these emissions some smoothing was applied to the transition between symbols, but to avoid more 'dead time'

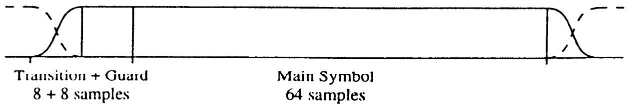

Figure 1: Transmitted OFDM symbol structure

it was incorporated into the guard period. The first 8 samples of the guard period are thus devoted to the transition between symbols, shaping the signal with a raised cosine. Figure 1 shows the symbol structure.

Due to the hardware used in the implementation some of these figures had to be changed slightly, as detailed in the next section.

Using some of the guard period for the symbol transitions may seem to defeat the object of having it, but in terms of the susceptibility to ISI it should make little difference here. The first half of the transition (ramping up to the new symbol and down from the previous one) contains more energy from the previous symbol than the current one, and as such could cause severe ISI. However, the excess delay required to shift the start of the guard period into the actual FFT window would typically be associated with very little signal energy (for an exponential channel), so the effect will be negligible. The second half of the transition contains progressively more energy from the current symbol and less from the previous one, which improves the ISI situation. The net effect of this is to reduce the guard period to approximately 12 samples (750 ns), but in the light of the original generous assumption of wanting 10 × the RMS delay spread, this still permits some 75 ns.

III Implementation

The baseband processing part of the transmitter uses an HP 8791 Model 7 Frequency Agile Signal Simulator (FASS), which generates the modulation as a bandpass signal. A two stage upconversion translates this first to around 1.5 GHz and then to 5.25 GHz, Figure 2. As a consequence of using the FASS there are certain restrictions on the use of sample frequency. It has a master clock of 2^{27} Hz, which when divided by 8 (due to oversampling everything by a factor of 8) gives 16.777216 MHz. A 64 point FFT running at this rate gives a symbol period of 3.8 μs, which is slightly less than the 4 μs listed previously and which increases the system data rate to 26.2144 Mbits/s. The guard period now needs more than the original 16 samples to maintain its duration of at least 1 μs and when coupled with the constraint of keeping the data throughput at more than 20 Mbits/s there are three possible sizes: 17, 18 or 19 samples. The middle value was chosen, which gives 78% efficiency and a throughput of 20.5 Mbits/s.

An OFDM signal can be differentially encoded in one of two ways: in the time domain it is a per-carrier method, where the phase of each individual carrier in the OFDM symbol is set by comparison with the phase on the same carrier from the previous symbol; in the frequency domain the encoding is done on a per-OFDM symbol basis and sets the phase of each carrier in the symbol with reference to the phase on the adjacent carrier in the same symbol. The time domain approach was used here.

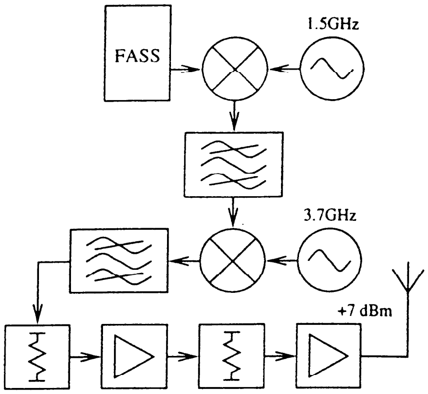

Figure 2: Transmitter block diagram

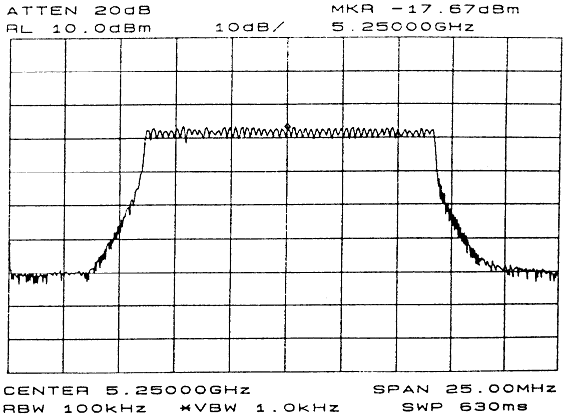

Figure 3: Transmitted spectrum

A continuous loop of 96 OFDM symbols separated by a 2 μs gap was transmitted from an omni-directional antenna at approximately +7dBm. The spectrum is shown in Figure 3, from which it can be seen that the signal bandwidth is approximately 15 MHz. This compares well with the 13 MHz Nyquist bandwidth of a 26 Mbits/s QPSK transmission, illustrating another positive feature of OFDM – its bandwidth efficiency.

The receiver follows a conventional design, first down-converting to 1.7 GHz, then to 70 MHz and finally in quadrature to baseband where it is sampled, Figure 4. Due to the oscillators and signal sources being accurate, high stability pieces of test equipment, no frequency control

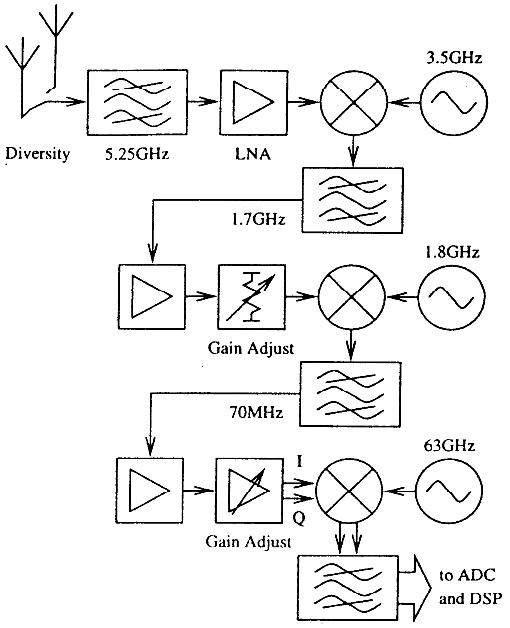

Figure 4: Receiver block diagram

loop is needed. There is also no carrier recovery because with differential decoding a precisely known carrier phase is not necessary. Gain control is in two parts: a switchable attenuator provides a coarse setting and a closed-loop variable gain amplifier just before the final stage provides the fine control. Synchronisation is also in two parts, comprising an energy detection algorithm to locate the gap between packets and then a correlation on the start of the packet to adjust this initial estimate.

Each channel (I & Q) is sampled by an 8 bit ADC, the outputs from which are then combined into a 16 bit word and stored in an 8 K × 18 bit FIFO. Once it is full, the captured signal is transfered to the DSP section of the receiver for processing and demodulating, before capturing the next block. Thus the system operates in a notional packet mode, not as a continuous 20 Mbits/s link, where the FIFO roughly corresponds to the size of a packet.

Because an OFDM signal has a large dynamic range but a low probability of actually reaching the peak values, accommodating all of it within the span of the ADC means that most of the time the signal is being sampled at less than the maximum resolution. The variable gain amplifier was therefore set to produce an output signal that would be clipped by the ADC when peaks occurred, but which would be sampled most of the time without clipping and with increased resolution. Simulations have shown that there is a negligible impact on the system error rate if the signal is clipped to a level of at least 3 × the RMS value; the amplifier here was set to deliver 4.5 × the RMS value.

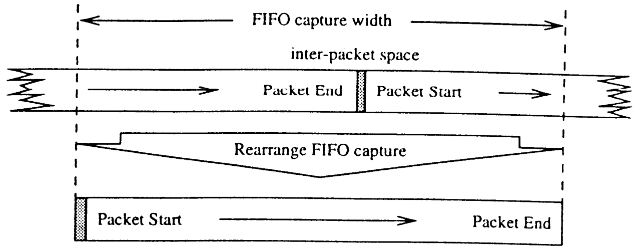

Figure 5: "Re-phasing" the captured packet

The DSP part of the system comprises two TIMs on a PC-mounted motherboard. The first of these is a TMS320C40, which is responsible for the FIFO interface and synchronisation, including reordering the samples to put the packet into the correct "phase", Figure 5. The samples are then transfered to the second TIM, which is a specialised FFT module from Kane Computing, comprising both a TMS320C44 and a BDSP9124 FFT processor running under its control. Each OFDM symbol in the packet is stripped of its guard period and passed through the BDSP9124 before differentially demodulating the signal and comparing it with a stored reference data pattern to determine errors.

IV Experimental Method

The transmitter was positioned in the laboratory such that it overlooked the office space of the building from an antenna height of 1.78 m, which is above the tops of the cubicle partitions. The receiver was assembled on a trolley with a pair of antennas between the top two shelves, corresponding to somewhere between desk height and the tops of the partitions but down amongst the shadowing and clutter of the office environment. (See photographs in Figures 9 and 10 at the end of the paper.) Alternate packets were processed from the different antennas, finding the maximum number of carriers in each OFDM symbol to produce an error in each case. The lower value from a consecutive pair of packets was kept as representing the output from an ideal two branch switched diversity scheme. The error statistics from 1000 diversity selected packets (*i.e.* 2000 transmitted packets) were collected from a range of locations chosen such that they lay on rings of 5 m, 10 m, 15 m and 20 m radius, concentric with the transmitter (see the plan at the end, Figure 11).

V Results and Interpretation

The first result to be presented is a calibration of the receiver sensitivity. This is shown in Figure 6. It should be noted that the filtering in the system was not optimised and approximately 3dB could probably be gained by rigorously matching the filter specifications to their respective tasks. However, their differential group delay will contribute to the overall ISI encountered by the signal and this will increase with tighter filters; there is already ap-

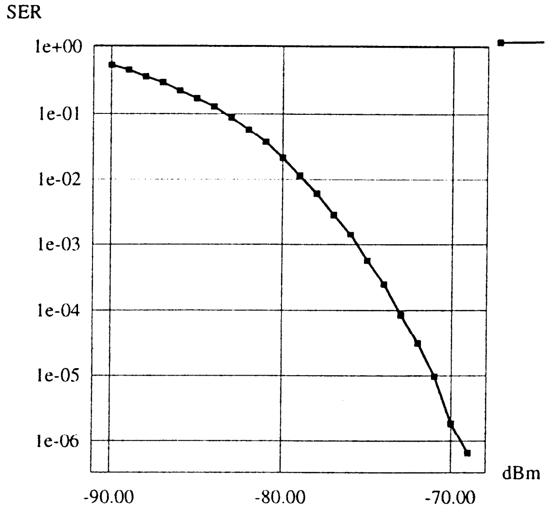

Figure 6: Receiver sensitivity: QPSK symbol error rate measured with transmitter and receiver back-to-back

proximately 30 ns of extra RMS delay spread from the filters used in this demonstrator. As OFDM is very sensitive to non-linearities, it is possible that improved performance could also be achieved by improving the linearity of some of the stages.

Figure 7 shows the spectrum of a typical signal received over a non-line-of-sight channel in the laboratory. The deep notch in the signal will clearly cause a much worse signal-to-noise ratio on the particular carrier affected, probably causing an error, even though the rest of the signal has been received in apparently good form. This illustrates how important coding is in OFDM: if the data on all the carriers is encoded with some redundancy, then errors arising from frequency selective fades can be corrected. The degree of error correction in the code is therefore an important system parameter which can be inferred from the results of these experiments.

Plotting the statistics of the number of carriers in error from each packet as a cumulative distribution produces curves such as those in Figure 8, which came from the 10 m ring. Each trace represents the results from a particular measurement location. From this Figure it can be seen that for 99% of the time there are less than 4 carriers in error from each OFDM symbol, across all the locations measured. Table 1 shows this, rounded up to the next integer, along with the results from the measurements at the other radii.

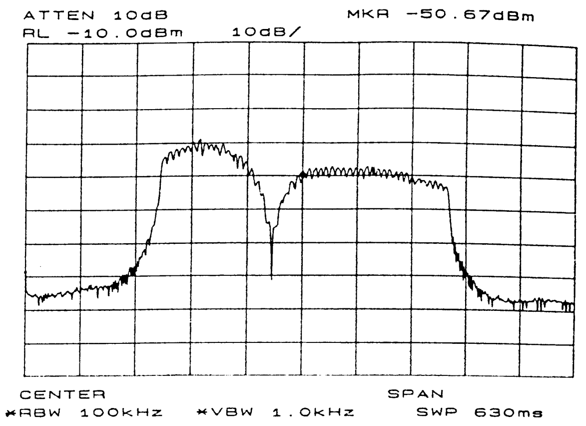

Figure 7: Received faded spectrum, measured at the LNA output

Radius	Carriers in Error
5	3
10	4
15	5
20	6

Table 1: Maximum number of carriers in error for 99% of the time at given distances from the transmitter

If the data in an OFDM symbol is encoded such that N carriers in error can be restored, then as long as the maximum number of carriers in error is less than N a packet will be received error free. The results in Table 1 showed how many carriers were in error for 99% of the time; correcting that number will therefore achieve a 99% success rate or alternatively a packet failure rate of 1%. Thus Table 1 provides an indication of the coding requirements of this OFDM system and environment.

VI Conclusions

A 20 Mbits/s OFDM demonstrator has been built and used to gather packet error rate statistics in an indoor office/laboratory environment. From these results the necessity of coding in OFDM has been amply illustrated, but more importantly the degree of correctability required of any code proposed for an OFDM system operating at a 1% packet failure rate has also been shown. As this correctability depends on added redundancy, the user data throughput is directly affected by the choice of code. This is therefore a useful contribution to system design.

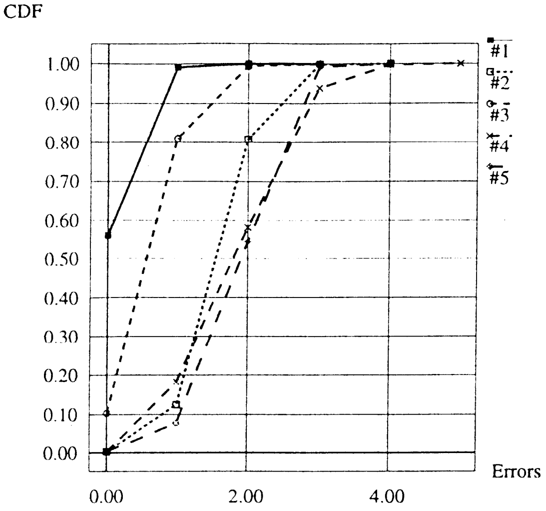

Figure 8: Cumulative distributions of the maximum number of carriers in error per OFDM symbol, per packet, for the 10 m ring

Figure 9: View over the measurement environment

References

[1] Technology Review Document, ETSI EP BRAN, wg3td38

[2] Nobles P. and Halsall F. "High Bit Rate Data Transmission within Buildings for Wireless Digital Cameras," IEE International Broadcasting Convention, September 1995.

[3] Hafezi P, Wedge D, Beach M.A. and Lawton M. "Propagation Measurements at 5.2 GHz in Commercial and Domestic Environments," IEEE PIMRC, September 1997.

Figure 10: View of the receiver on its trolley

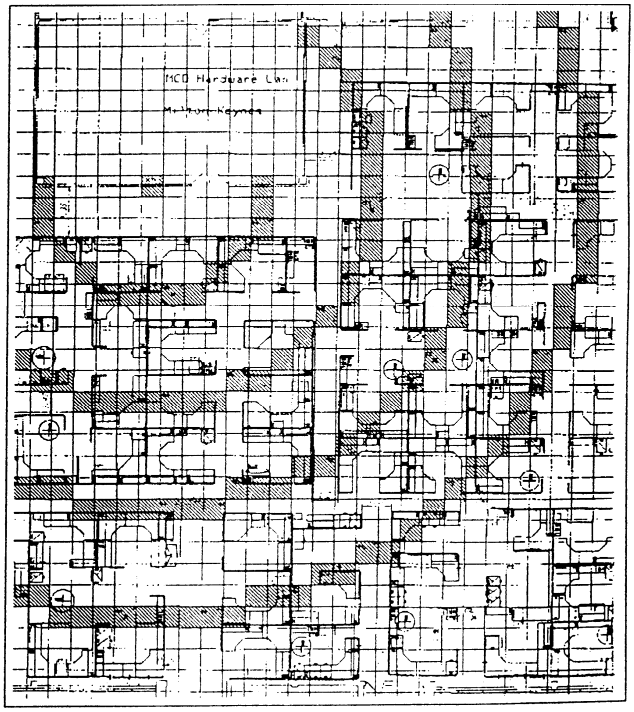

Figure 11: Plan of the measurement environment with a 1 m grid