
An Unconventional, Highly Multipath-Resistant, Modulation Scheme

Study Leader:
W. Press

Contributors Include:
W. Dally
D. Eardley
R. Garwin
P. Horowitz

September 1997

JSR-97-160

DTIC QUALITY INSPECTED 4

Approved for public release; distribution unlimited.

JASON
The MITRE Corporation
1820 Dolley Madison Boulevard
McLean, Virginia 22102-3481
(703) 883-6997

Contents

| | | |
|----------|---|-----------|
| 1 | INTRODUCTION | 1 |
| 2 | SLOW BIT, BINARY ON/OFF KEYING (BO/OK) | 5 |
| 2.1 | Bit Error Rates for BPSK versus BO/OK | 5 |
| 2.2 | Channel Capacity for BPSK versus BO/OK | 10 |
| 3 | THE USES OF DIVERSITY | 13 |
| 4 | CODING FOR A HIGH-BER ASYMMETRIC CHANNEL | 21 |
| 4.1 | “Existence Proof” | 21 |
| 4.2 | Sketch of More Efficient Coding Schemes | 22 |
| 5 | HYPERCARRIER PARALLELISM | 25 |
| 5.1 | Direct FFT Synthesis of Waveforms | 25 |
| 5.2 | Hardware Implementation | 27 |
| 5.2.1 | Power, Size, Cost | 30 |
| 5.2.2 | Processing | 31 |
| 5.2.3 | Variations | 32 |
| 5.3 | Extra Spreading via “Tukey’s Trick” | 32 |
| 5.3.1 | Implementation | 34 |
| 6 | Reducing the Detectability of Hypercarrier | 35 |
| 6.1 | Sparse Hypercarrier | 35 |
| 7 | SUMMARY | 37 |

1 INTRODUCTION

In the obstructed urban environment, the RF channel between two mobile communicators (whom we will here regard as being pedestrians on foot) is degraded in two distinguishable ways.

First, the *mean* path loss is greater than the ideal free-space inverse square law. There exist a variety of models for this, none very accurate because of the dependence on details of the environment. Some references characterize the urban path loss as a power law with exponent between about -3 and -5 (as compared to free-space -2), but this clearly must hold (if at all) over only a limited range of distance and conditions. Relevant physical effects include reflection, diffraction, scattering, and absorption. (See Rappaport, 1996.)

Second, because of reflections and shadows, the signal is subject to fading. Shadow fading, so called, occurs when the received signal is substantially blocked by objects. Multipath fading, so called, occurs due to the interference of multiple copies of the signal arriving at the receiver at different times. Destructive interference can cause the received signal power to decrease by a large factor from its mean value. (A quantitative treatment is below.)

Path loss and shadow fading are approximately independent of frequency, at least over any restricted bandwidth $\Delta f/f \ll 1$. There is thus essentially no counter to them except to increase transmitted signal power or (for shadow fading) to retransmit lost information when the transmitter or receiver emerges from shadow. Equivalently, a coding technique with forward error correction and very long interleaving (such as the now-in-vogue "Turbo codes") could be used, but not always feasibly for high data rate transmissions.

Note that operating at relatively low (VHF) frequencies reduces shadow fading somewhat as the signal will diffract around obstacles in the urban environment.

Multipath fading is more insidious and interesting. Because it is an interference effect, the fading pattern of a given frequency can vary on a spatial scale on the order of a wavelength. Uncorrected bit error rates will therefore be high, as the communicators move through fades of various depths.

Also, different frequencies, even very close, can have uncorrelated fading patterns. A consequence is differential fading across the information sidebands of a conventionally modulated signal, which produces intersymbol interference. In order of magnitude, this occurs when the product of the signal bandwidth and the RMS delay spread is greater than 1. Simple, conventional modulation schemes require, in practice, a value less than 0.2. Since the RMS delay spread in the urban environment can be as large as $25 \mu\text{s}$, we see that data rates greater than 8 kilobits per second (kbps) require specialized modulation schemes even when the instantaneous received power is seemingly adequate.

In the commercial world, modulation schemes that are *somewhat* resistant to multipath fading are combined with increased signal power, and with careful choices for the locations of at least one end of the channel (the cellular base station), to achieve workable systems. Furthermore, up to now, the data rates of interest commercially have been fairly low, suitable for voice cellular services.

For the problems that we are here interested in, the commercial solutions are generally not workable: (1) Our transmitted power must be low, both for battery life and LPI/LPD reasons. (2) Both ends of our link are at "random" places in the environment. (3) We would like to achieve much higher data rates, at least in an occasional burst mode.

In some applications, multipath effects can be overcome by the use of a RAKE receiver, which combines *coherently* the multiple independent copies of the transmitted signal that it is able to detect. This combination eliminates much of the inter-symbol interference that results when the variation in path length is greater than the length of a symbol. It is not, however, able to overcome fading since the same fading signal is input to all taps or 'fingers' of the RAKE.

We are therefore left with the thought that it would be good to have a modulation scheme with the following properties:

- Capable of high (e.g., 1 Mbps) data rates
- Completely resistant to multipath effects, including both multipath fading and intersymbol interference
- Implementable in a small package in current technology
- Good LPI/LPD properties (to be discussed below)

Surprisingly, we think that such a modulation scheme does exist, and we describe it in the remainder of this paper. Although our study is purely "paper", we think that a good engineering design team (including good expertise in digital coding) could field a functioning prototype system for field testing.

It is important to say what our modulation scheme does *not* do!

It does not

- Solve the problem of (mean) path loss in the urban environment

Therefore, we will not, as the reader might otherwise wish, compute the transmitter power necessary for communication as a function of bit rate

and distance. What we *will* do is compare quantitatively the performance of the modulation scheme here proposed to the performance of a standard modulation scheme (BPSK) for the same peak power transmitted and same path loss.

In other words, the user should imagine a fiducial BPSK system operating with the urban environment's path loss, but (magically) with *no* multipath effects. We will give the performance of our proposed scheme – *with* multipath effects – relative to that fiducial BPSK system. These results will enable an actual prediction for our scheme's performance (as a function of bit rate, distance, and transmitted power) to be easily calculated, once one is given actual measurements of mean received power in an actual environment. (Such measurements are not difficult.)

2 SLOW BIT, BINARY ON/OFF KEYING (BO/OK)

We will discuss a series of concepts, none new individually and none seemingly promising for high data-rate links, but which when taken together will make a workable system. The first concept is slow binary off-on keying (BO/OK).

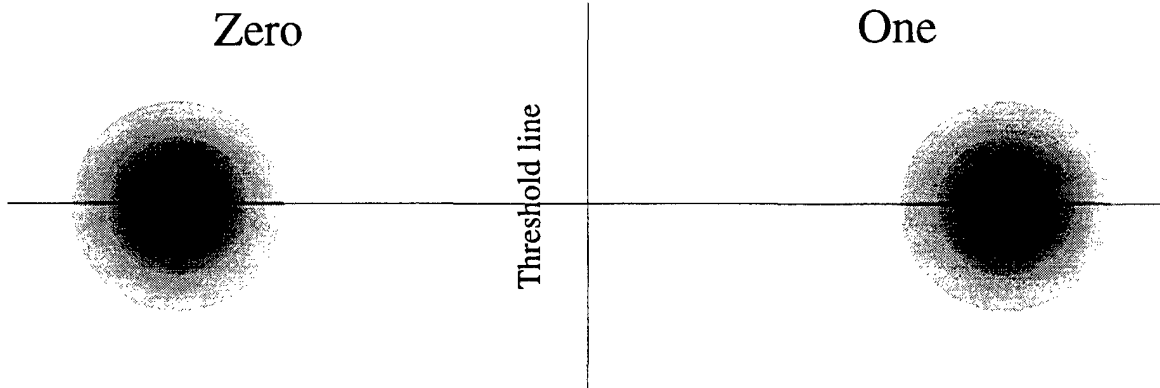
Suppose we send a bit by keying a pure carrier for a time long compared with the “ringing time” or RMS delay spread of the environment ($100 \mu\text{s}$ is a good value to keep in mind). A “1” value is sent as carrier on, while a “0” value is sent as carrier off, just like ancient CW Morse. Taken at face value, this method is both “slow” (maximum of 10^4 bps with the above parameters) and, as we will now see, error-prone (high error rate).

At the receiver, multiple copies of the carrier combine incoherently. In the worst-case limit of a large number of copies of comparable amplitudes, there is no predictable phase relationship between the transmitted and received signals. In fact, the received signal has the statistics of a Rayleigh fading channel, whose most probable received amplitude, in the phasor (complex amplitude) plane, is zero. This sounds terrible, but let us persist and compare this channel to the case of BPSK.

2.1 Bit Error Rates for BPSK versus BO/OK

Figure 1 shows schematically the phasor plane for both BPSK and BO/OK. In BPSK, bits are transmitted as positive or negative values on the real axis, with magnitudes set by the peak power. They are received, with Gaussian noise added, as Gaussian probability distributions centered on the

BPSK



BO/OK

(Rayleigh limiting case of infinite multipath)

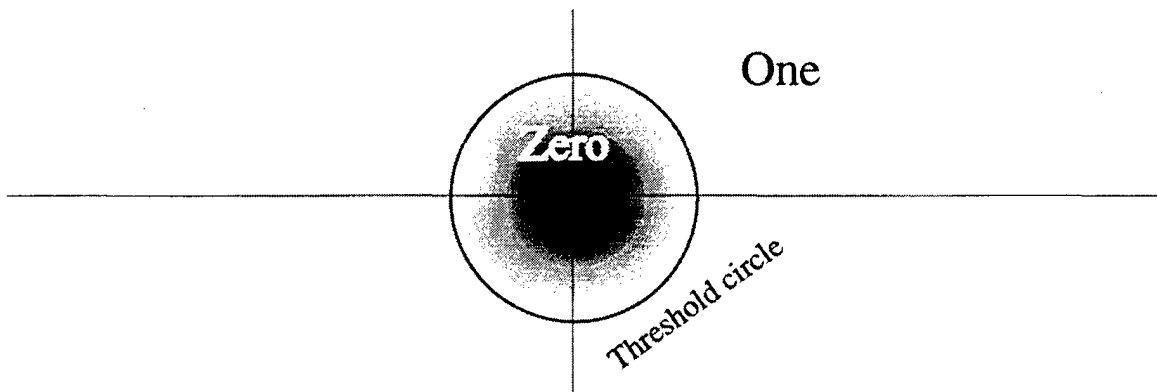


Figure 1. Phasor plane comparison of BPSK and BO/OK modulation schemes BPSK is shown without any multipath fading. BO/OK is shown in the limit of total (Rayleigh) multipath fading.

transmitted points. The receiver distinguishes zero from one by the sign of the real part of the received signal. By symmetry, this must be the optimal threshold. However, more generally, we might imagine setting the dividing line at some other position on the real axis. If σ is the one-dimensional noise standard deviation in units of the signal strength (thus, normally, $\sigma \ll 1$), so that

$$1/\sigma^2 = S/N = \text{Signal-to-Noise Ratio (Power)} \quad (1)$$

and we measure the position t of the threshold also in units of the signal amplitude, then the bit error rate (BER) for a transmitted zero is

$$\text{BER} = \frac{1}{2} \text{erfc} \left(\frac{1+t}{\sqrt{2}\sigma} \right) \quad (2)$$

while the BER for a transmitted one is

$$\text{BER} = \frac{1}{2} \text{erfc} \left(\frac{1-t}{\sqrt{2}\sigma} \right) \quad (3)$$

These relationships are shown schematically (i.e., for a particular value of S/N) in the top half of Figure 2. One sees that as a general property of the error function “erfc” there is a broad “window” of thresholds (including the symmetrical value 0) where the bit error rates are exponentially small for both transmitted zeros and ones.

Now compare the case of binary off/on keying in a Rayleigh fading environment, as shown in the bottom halves of Figures 1 and 2. When we send a zero, it is received as a Gaussian distribution of one-dimensional standard deviation σ , as in the BPSK case.

When we send a one, however, Rayleigh statistics cause it to be received as a fat Gaussian of standard deviation $\sqrt{1+\sigma^2}$ (in these units, the signal contributes the 1) *also centered on the origin*. In this case, the best we can do is to threshold on complex amplitude, shown as a circle in Figure 1, since all phase information is (presumed) lost. This is simply pure “presence-of-tone” detection.

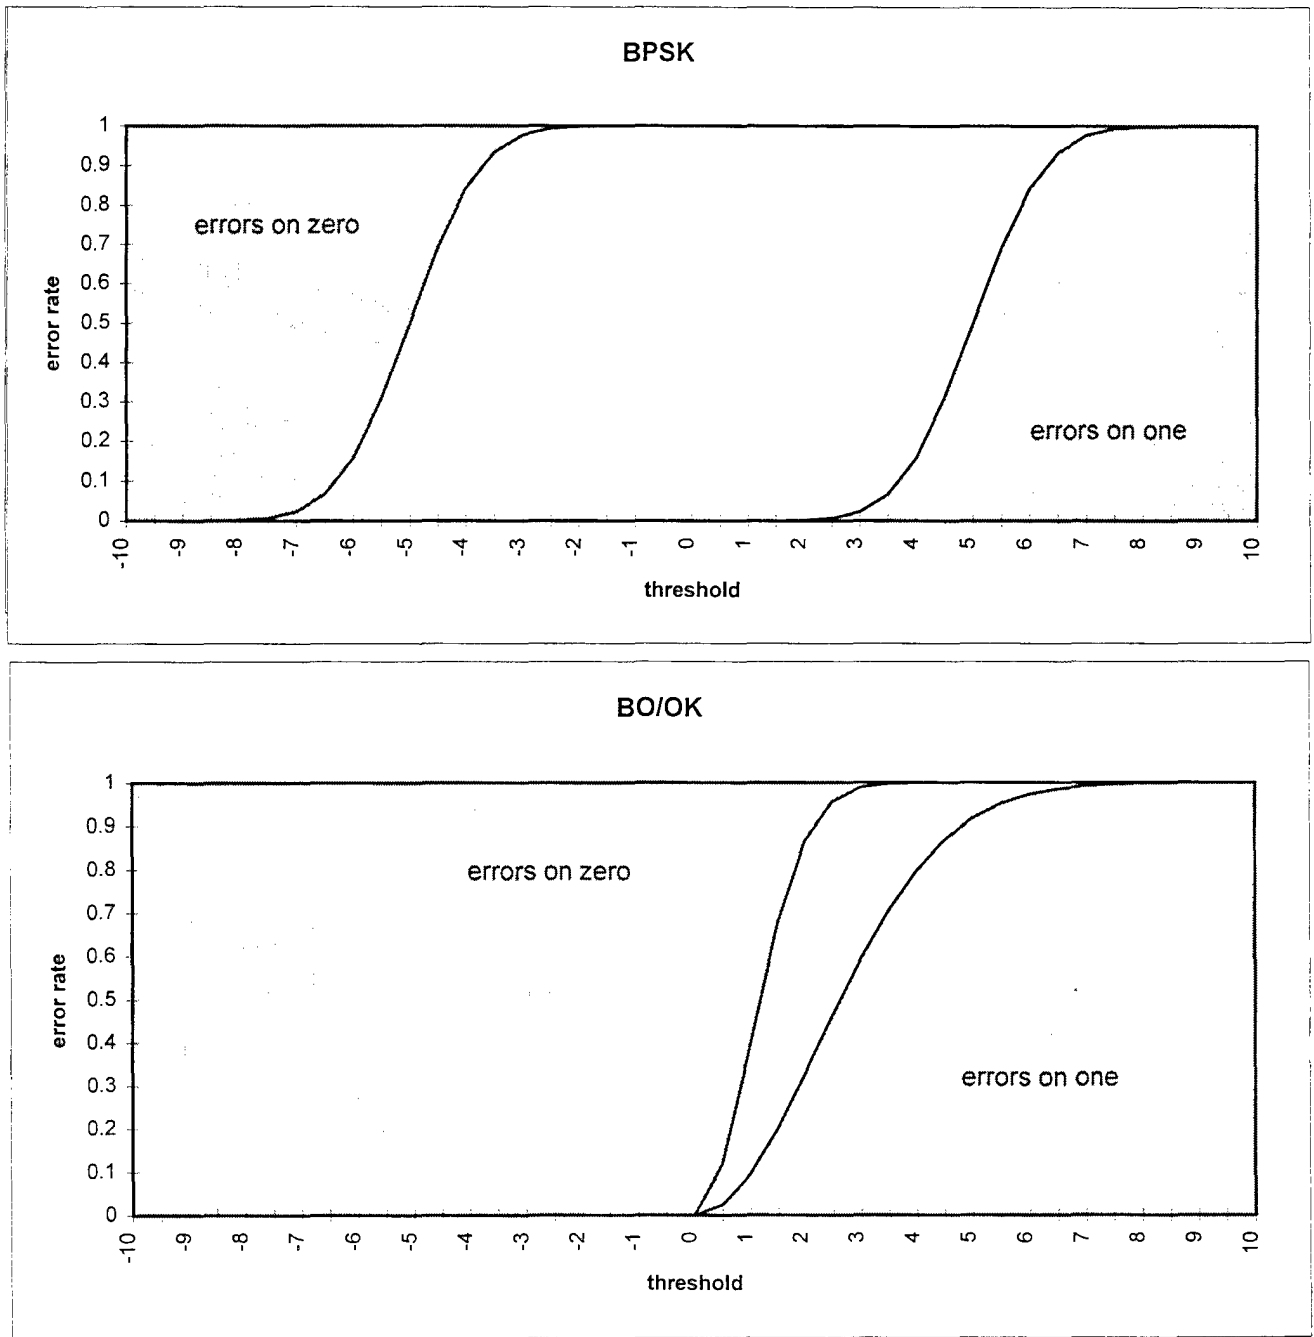


Figure 2. Bit Error Rates for BPSK and BO/PK as a function of the threshold (line and circle as shown in Figure 1).

The error rates for zero and one bits are now quite unsymmetrical, and depend on the radius chosen for the threshold amplitude. The probability distribution for the complex modulus amplitude $|a|$ with tone absent is

$$p(|a|)da \propto \frac{|a|}{\sigma^2} \exp\left(-\frac{|a|^2}{2\sigma^2}\right) \quad (4)$$

With tone present, it is

$$p(|a|)da \propto \frac{|a|}{(1 + \sigma^2)} \exp\left(-\frac{|a|^2}{2(1 + \sigma^2)}\right) \quad (5)$$

Integrating this and taking a threshold radius t , the BER for a transmitted zero is

$$\text{BER} = \exp\left(-\frac{t^2}{2\sigma^2}\right) \quad (6)$$

which goes to zero exponentially as t becomes large, while the BER for a transmitted one is

$$\text{BER} = 1 - \exp\left(-\frac{t^2}{2(1 + \sigma^2)}\right) \quad (7)$$

which goes to zero only polynomially as t goes to zero.

The behavior of these BER expressions is shown schematically in the bottom half of Figure 2. One sees that, by contrast with BPSK, there is no large "open window" for choice of threshold.

We describe in the next subsection how to pick an optimal value for the threshold t , but some indicative numbers may be helpful here: When $S/N = 10$ (that is, $\sigma = 0.316$), the optimal threshold turns out to be $t = 0.82$ (again, in units of the carrier-on amplitude), and the probabilities of receiving correct or incorrect bits (transition matrix components) are

| | | | | |
|---------|---|------------|-------|--|
| | | Receive as | | |
| | | 0 | 1 | |
| Send as | 0 | 0.965 | 0.035 | |
| | 1 | 0.263 | 0.737 | |

(8)

So ones are received as zeros 26% of the time! Note, however, that BO/OK is *completely* insensitive to multipath effects. Indeed, the assumption

of a Rayleigh distribution is equivalent to assuming the most extreme case of many paths of comparable amplitude. Any less extreme case will result in smaller BERs for transmitted “one” bits, because the carrier-on amplitude distribution will be more peaked away from the origin.

By contrast, the above analysis for BPSK assumes *no* multipath degradation. BPSK (without special fixes) is in fact known to degrade badly in a multipath environment.

So, at this stage we have a seemingly poor, but definitely multipath insensitive, modulation scheme.

2.2 Channel Capacity for BPSK versus BO/OK

A surprise comes when we compute the theoretical channel capacity of a “binary asymmetric channel” (BAC) such as the one above. We will find that, for our typical parameters, it is only a factor of two or so less than a BPSK channel!

The entropy function $H(\mathbf{x})$ as a function of a vector of probabilities \mathbf{x} (whose components sum to 1) is defined by

$$H(\mathbf{x}) \equiv - \sum_i x_i \log_2 x_i \quad (9)$$

The mutual information between two random variables S and R (for “send” and “receive”) is defined by

$$I(S, R) = H(\mathbf{p}_S) + H(\mathbf{p}_R) - H(\mathbf{p}_{S \otimes R}) \quad (10)$$

If, now, S represents the probabilities with which zeros or ones are sent, R the probabilities with which they are received, and $S \otimes R$ represents the

2×2 contingency table of a zero or one being sent or received, with its 4 values summing to 1, then the capacity of the channel is given by

$$C = \max_{\mathbf{p}_S} I(S, R) \quad (11)$$

The reason for the “max” is that we get to choose the probability with which we send zeros or ones, that is, \mathbf{p}_S . (A pre-coding step would transform any input distribution of zeros and ones to the desired probability.) Channel capacity is measured in “bits per bit”, that is, decoded error-free bits per actual bit sent.

We now can say how the threshold t is optimized for BO/OK: It is chosen to maximize the channel capacity. While in principle the channel capacity involves maximizing over \mathbf{p}_S , the probability with which we send zeros and ones, in practice, we have found that this extra maximization gives only a very small increase in channel capacity. Therefore, in what follows, we always take $\mathbf{p}_S = (0.5, 0.5)$.

The following table gives, for different S/N ratios, the optimal threshold t , channel capacity C , and the four components of the transition matrix. (The final column is explained in Section 3, below.)

| S/N(dB) | t | C | transition matrix | | | | 1-BER (1/6 code) |
|---------|------|-------|-------------------|-------|-------|-------|------------------|
| -10 | 5.76 | 0.001 | 0.809 | 0.191 | 0.778 | 0.222 | 0.538 |
| -8 | 4.65 | 0.003 | 0.819 | 0.181 | 0.772 | 0.228 | 0.560 |
| -6 | 3.75 | 0.006 | 0.829 | 0.171 | 0.756 | 0.244 | 0.591 |
| -4 | 3.06 | 0.013 | 0.845 | 0.155 | 0.736 | 0.264 | 0.633 |
| -2 | 2.51 | 0.027 | 0.862 | 0.138 | 0.703 | 0.297 | 0.689 |
| 0 | 2.06 | 0.053 | 0.880 | 0.120 | 0.654 | 0.346 | 0.760 |
| 2 | 1.71 | 0.097 | 0.901 | 0.099 | 0.591 | 0.409 | 0.834 |
| 4 | 1.42 | 0.159 | 0.920 | 0.080 | 0.514 | 0.486 | 0.900 |
| 6 | 1.18 | 0.241 | 0.937 | 0.063 | 0.426 | 0.574 | 0.948 |
| 8 | 0.98 | 0.336 | 0.953 | 0.047 | 0.341 | 0.659 | 0.975 |
| 10 | 0.82 | 0.437 | 0.965 | 0.035 | 0.263 | 0.737 | 0.989 |
| 12 | 0.68 | 0.537 | 0.975 | 0.025 | 0.196 | 0.804 | 0.995 |
| 14 | 0.57 | 0.628 | 0.982 | 0.018 | 0.143 | 0.857 | 0.997 |
| 16 | 0.47 | 0.708 | 0.988 | 0.012 | 0.102 | 0.898 | 0.999 |
| 18 | 0.39 | 0.776 | 0.992 | 0.008 | 0.072 | 0.928 | 1.000 |
| 20 | 0.32 | 0.830 | 0.994 | 0.006 | 0.050 | 0.950 | 1.000 |
| 22 | 0.27 | 0.873 | 0.996 | 0.004 | 0.035 | 0.965 | 1.000 |
| 24 | 0.22 | 0.907 | 0.998 | 0.002 | 0.024 | 0.976 | 1.000 |
| 26 | 0.18 | 0.932 | 0.998 | 0.002 | 0.016 | 0.984 | 1.000 |
| 28 | 0.15 | 0.951 | 0.999 | 0.001 | 0.011 | 0.989 | 1.000 |
| 30 | 0.12 | 0.965 | 0.999 | 0.001 | 0.007 | 0.993 | 1.000 |

(12)

Figure 3 shows channel capacity C as a function of S/N for both BO/OK (the curve labeled “1”) and for BPSK (so labeled). The other curves, not further discussed here, show the improvement that would derive from antenna diversity with 2, 3, 4, and 10 antennas. Multiple antennas act to decrease the amount of fading of the Rayleigh channel.

We see that for $S/N > 10\text{dB}$ (a likely range in which we would want to operate), the channel capacity ranges from 0.437 to nearly 1, despite the relatively large BER values (especially for transmitted “ones”). This means that, with good forward error correcting coding, BO/OK is a feasible – and completely multipath insensitive – modulation scheme. In Section 4 we will discuss how such coding can be achieved. But first we will discuss the use of antenna diversity to ameliorate channel fading.

3 THE USES OF DIVERSITY

In a multipath environment, diversity of propagation path can often be an effective cure for fading. The link can use more than one receiving antenna, or more than one transmitting antenna, or both. For discussion see J.G. Proakis, *Digital Communications*, (New York: McGraw-Hill, 1982), Chapter 7. In this section we shall study the bit error rate for BO/OK as it depends on path diversity.

We assume that some number n of propagation paths are utilized, different enough that the statistics of signal strength, noise strength, and fading are independent for different paths. For instance, we might use n receiving antennas spaced far enough apart to receive the signal from a single transmitting antenna. According to lore, antenna spacing should be 10λ or greater for long-distance HF or tropo scatter systems; however urban multipath environments seem likely to be different, and antenna spacings $\sim \lambda$ may well suffice. A prosaic example is FM reception in a car; nulls due to multipath fading are often noticed to be separated by $\sim \lambda/2$.

Our strategy for combining the signals from the different antennas is simple: We sum the detected powers; *i.e.*, take the r.m.s. of the detected amplitudes. It seems likely this is the optimal strategy but we have not proved so.

The probability distribution for the amplitude $|a|$ of the summed signal is

$$p_n^{\text{signal}}(|a|) \propto \frac{|a|^{2n+1}}{(1 + \sigma^2)^n} \exp\left(\frac{-|a|^2}{2(1 + \sigma^2)}\right) \quad (1)$$

while the probability distribution for the noise amplitude in the summed signal is

$$p_n^{\text{noise}}(|a|) \propto \frac{|a|^{2n+1}}{\sigma^{2n}} \exp\left(\frac{-|a|^2}{2\sigma^2}\right) \quad (2)$$

Taking a signal threshold of t for detection of a 1-bit, the BER for a transmitted zero is

$$\text{BER} = p_n \left(\frac{t^2}{2\sigma^2} \right) \exp \left(-\frac{t^2}{2\sigma^2} \right) \quad (3)$$

while the BER for a transmitted one is

$$\text{BER} = 1 - f_n \left(\frac{t^2}{2(1 + \sigma^2)} \right) \exp \left(-\frac{t^2}{2(1 + \sigma^2)} \right) \quad (4)$$

where the polynomial $f_n()$ is defined as

$$f_n(x) = \sum_{k=1}^{n-1} \frac{x^k}{k!} \quad (5)$$

Note that $f_n(x)$ is a truncation of the Taylor series for $\exp(x)$ and is a good approximation for $x \lesssim n$, so that the probability distributions $p_n()$ are close to 0 or 1 over most of their range with a sharp transition at argument $x \sim n$.

Figures 3–7 compare the performance of BO/OK at various path diversities $n = 1, 2, 3, 4, 10$ in a strong multipath environment, and also compare these with the performance of conventional BPSK (binary phase shift keying) with *no* multipath degradation and a single path ($n = 1$), all at the same S/N ratio. The main conclusion is shown in Figure 3, which shows that channel capacity increases with diversity, as expected; in fact $n = 10$ BO/OK outperforms $n = 1$ BPSK. These results for channel capacity presume that the detection threshold t has been optimized (Figure 5), and that the optimal proportion of 1-bits is utilized in the transmitted signal (Figure 4). (With respect to the latter, the calculations of this section differ from those of Section 2.1 where the proportion of transmitted 1-bits was fixed at 0.5; however this makes little difference in the channel capacity.)

The bit error rates p (for a transmitted 1-bit) and q (for a transmitted 0-bit) are displayed in Figures 6 and 7.




Figure 3. Theoretical channel capacity of BPSK (without multipath) and BO/OK with multipath) as a function of Signal-to-Noise power ratio. The labeled parameter shows the number of assumed antennas in the case of diversity.




Figure 4. Theoretical optimal density of transmitted 1-bits for BPSK modulation in the presence of Rayleigh fading.




Figure 5. Optimal receiver threshold amplitude for BO/OK as a function of Signal-to-Noise power ratio




Figure 6. Bit Error Rate for transmitted 1-bits as a function of Signal-to-Noise power ratio.




Figure 7. Bit Error Rate for transmitted 0-bits as a function of Signal-to-Noise power ratio.

4 CODING FOR A HIGH-BER ASYMMETRIC CHANNEL

4.1 “Existence Proof”

The previous calculations of theoretical channel capacity show that an optimal 1/2 rate code (two transmitted bits per information bit) should be more than enough to achieve arbitrary small system BERs (for an assumed $S/N > 10\text{dB}$). However, most well-known coding techniques cannot be directly applied to the large raw BERs that we have.

It is therefore useful to demonstrate explicitly that a simple 1/6 rate code (which throws away an additional factor of 3 in rate) is adequate to bring our BERs into the regime where standard codes apply.

We take the simplest possible code: A zero bit is sent as 6 zeros; a one bit is sent as 6 ones. You may want to think of these not as consecutive bits, but as interleaved over a significant time, so that they become statistically independent in an environment of time-dependent fading. (In Section 4 we will discuss a quite different method.)

As for decoding, a short calculation (or simulation) is required to find the optimal decode. It is not optimal to go by majority vote of zeros or ones, e.g., because the channel is asymmetric. The answer (again for $S/N = 10\text{dB}$) turns out to be that patterns with 6 or 5 zeros should decode to zero, while patterns with two or more ones should decode to one.

In the table in Section 2.2, above, the last column gives the BER achieved using only this trivial code. One sees that a BER of 10^{-2} is achieved

for $S/N = 10\text{dB}$. While this is not adequate as an overall system performance, it is small enough that an additional layer of standard coding – with very little overhead in terms of signal rate – can achieve any desired system BER.

We will see below that even with the extra factor of 3 hit taken, megabit per second data rates are achievable.

4.2 Sketch of More Efficient Coding Schemes

Using an efficient forward error correcting code, one can more closely approach the theoretical capacity of the BO/OK channel. Because the modulation scheme is itself multipath insensitive, the coding scheme need not be designed to deal with the long burst errors typical of fading channels. While the raw bit error rate is high, the error statistics of adjacent bits are to first order uncorrelated. Thus the complexity of using a long interleave or a concatenated code (such as a *turbo code*) to achieve a long constraint length is not required.

The characteristics of the BO/OK channel suggest the use of a relatively short constraint-length convolutional code used in conjunction with a *soft decision* trellis decoder. Consider, for example, the rate 1/3 convolutional code with constraint-length 5 generated by the polynomials:

$$x^4 + x^2 + 1 \quad (1)$$

$$x^4 + x^3 + x + 1 \quad (2)$$

$$x^4 + x^3 + x^2 + x + 1 \quad (3)$$

This code has a *free distance* of 12, and is thus able to correct all patterns of up to 5 errors over a 15-code bit interval. Assuming hard detection with a threshold set to give a symmetric channel (both pessimistic assumptions),

the probability correctly detecting all of the bits in a block of $d=15$ code bits (5 information bits) is bounded by

$$p_b > \sum_{m=0}^5 \binom{d}{m} p^{d-m} (1-p)^m \quad (4)$$

With a hard threshold of $t = 0.59$ the probability of correct detection is $p = 0.82$ for both a one and a zero which gives a bound of correct block detection of $p_b = 0.985$. This is a conservative bound as it does not take into account the large numbers of patterns of greater than 5 bit errors that are also corrected by this code. The calculation is also quite pessimistic as it does not take advantage of the increased capacity that comes from operating the channel using an asymmetric detection threshold or by using a soft decision decoder. Even so, the calculation shows that we can achieve acceptable bit error rates with a capacity that is within a factor of 1.5 of the theoretical capacity calculated above. A simulation of the channel operating with a soft-decision decoder should be conducted to more accurately quantify the actual bit-error rate.

The encoder keeps a history of the last five information bits in a shift register. Each time a new information bit is shifted in, three code bits are generated by *multiplying* the information bits (considered to be a polynomial over GF(2)) by the three generator polynomials.

The asymmetrical nature of the BO/OK channel is a natural match for a soft decision trellis decoder. With a soft-decision decoder, there is no need for a threshold, t . Instead, for each received symbol, the receiver detects the actual signal magnitude, a , and calculates $p_1(a)$ and $p_0(a)$, the probability that the symbol is a 1 or a 0 respectively, given that the received value is a . These probabilities are then used to propagate a search over a trellis of possible received codewords to determine the most likely codeword given the history of received amplitudes. For typical codes and channels, the use of soft decoding gives the equivalent of a 3dB increase in S/N.

Soft-decision decoding for the constraint-length 5 convolutional code described here can easily be performed in software using either a Viterbi decoder or a trellis decoder. A Viterbi decoder maintains 32 states (corresponding to the last 5 information bits) and for each state records the most likely path to have reached that state. For each received frame, the decoder extends each of these paths by one step using the received signal magnitudes. The transition probability from each state to each possible next state (eight in our example) is computed and the most likely path to reach each next state is retained. If all paths begin with the same symbol, then the decode is successful and the beginning symbol is shifted out of the decoder as the received bit.

The inner-loop of the Viterbi decoder computes 256 possible transitions each frame. At about 10 operations per transition, a total of about 2G operations/s is required to maintain a 1Mb/s data rate. This number can be reduced by an order of magnitude by using a trellis decoder that extends paths from only the most likely few states. Using a code with a shorter constraint length reduces the required number of operations exponentially.

5 HYPERCARRIER PARALLELISM

To recap: We have achieved a multipath insensitive modulation with (after forward error correction) arbitrarily small bit error rates. However, we are still limited to a raw (uncorrected) bit rate on the order of 10^4 bps, because of the requirement that each bit last for the ringing time of the urban environment.

To increase the signalling rate by 3 orders of magnitude, we will send and receive about 10^3 channels simultaneously, using direct waveform synthesis by FFT methods, and also using FFTs on the receive end. This is a form of “multicarrier modulation”. However, multicarrier modulation, as the term is usually used, refers to dividing the information bandwidth into a *small number* of separate channels, each narrow enough so as to avoid intersymbol interference (ISI) in the presence of an RMS delay spread (see Section 1, above). Here, by contrast, ISI is not an issue, because we are sending “slow bits”, perhaps lasting $100 \mu\text{s}$ each. The purpose of our multicarrier modulation is simply to achieve parallelism on these bits. Since our number of carriers will be much larger than the conventional case, we will use the term “hypercarrier” instead of multicarrier.

5.1 Direct FFT Synthesis of Waveforms

The basic architecture of a hypercarrier transmitter is: (1) digital processor (performing FFTs), to (2) D/A converter, to (3) upconverter (from baseband to the desired RF band), to (4) final amplifier. The basic architecture of the receiver is: (1) downconverter, to (2) A/D converter, to (3) digital processor. In other words, this is a fully digital radio. We will see that current commercial components are available to build practical hypercarrier

systems.

The design space for a hypercarrier system is set by the following parameters:

$$\tau = \text{duration of a bit } (\mu\text{s}) \quad (1)$$

$$N = \text{length of complex FFT} \quad (2)$$

$$r = \text{A/D sample rate (Hz) for complex sampling} \quad (3)$$

$$f_0 = \text{central frequency after upconversion (Hz)} \quad (4)$$

$$\text{BW} = \text{RF bandwidth (after upconversion) (Hz)} \quad (5)$$

In the next section we will add an additional parameter,

$$K = \text{number of repetitions of waveform per bit} \quad (6)$$

In this section we take $K = 1$. Assuming Nyquist critical sampling, we have the relation (for complex samples)

$$\text{BW} = r \quad (7)$$

In the simplest design, we perform one FFT on the sampled signal every time τ , so that we have

$$r\tau = N \quad (8)$$

A sample design might have $\tau = 100 \mu\text{s}$, $N = 1024$, $r = 10 \text{ MHz}$, $\text{BW} = 10 \text{ MHz}$. A weakness of this design is that the bandwidth is not as large as one would like in all cases, both to achieve optimal LPI/LPD, and also to have uncorrelated multipath fading across it. In some situations, however, this design may meet all requirements, so we give some further design hints here.

5.2 Hardware Implementation

Figure 8 shows a straightforward receiver implementation for a hyper-carrier digital communication system as just described. Here we have chosen COTS surface-mount components, to demonstrate that a compact receiver can be easily built to meet the requirements; in practice one might wish to invest in a custom MMIC implementation for smaller size and, likely, lower per-unit production cost. For VHF operation at $f_0 = 250$ MHz we are dealing with antenna sizes of order $L \approx \lambda/4 \approx 30$ cm; the receiver architecture is, however, independent of operating frequency.

The receiver is a “single conversion to complex baseband” design. Because the antenna sees a 300K (minimum) thermal noise environment, typically augmented by some 10’s of dB of “cultural” noise and interference (see Figure 5), there is essentially no benefit to using a low-noise amplifier at the front end; the inexpensive MAR-6SM (\$1.21), with 20 dB gain, 3 dB noise figure, and low-voltage operation (+3.5V, 16 mA) will do the job. Furthermore, it has dc–2 GHz bandwidth, so the design is easily adapted to higher operating frequencies. The RF bandpass filter (which might be incorporated into a tuned frontend amplifier) is a good idea to block out of band signals and thus preserve headroom in the second amplifier and mixer; the unit indicated is a compact LC design. The second RF amplifier, which can be implemented as a cascade of inexpensive MAR-type amplifiers, brings the signal level up to ≈ -20 dBm for mixing to baseband.

The mixer configuration (sometimes called a complex demodulator, and obtainable as a single component, e.g. the MCL JCIQ-series) uses an LO at RF midband to drive a pair of mixers through a quadrature splitter, producing a quadrature pair of baseband voltages whose spectrum extends from -5 MHz to $+5$ MHz – the translated 10 MHz bandpass originally

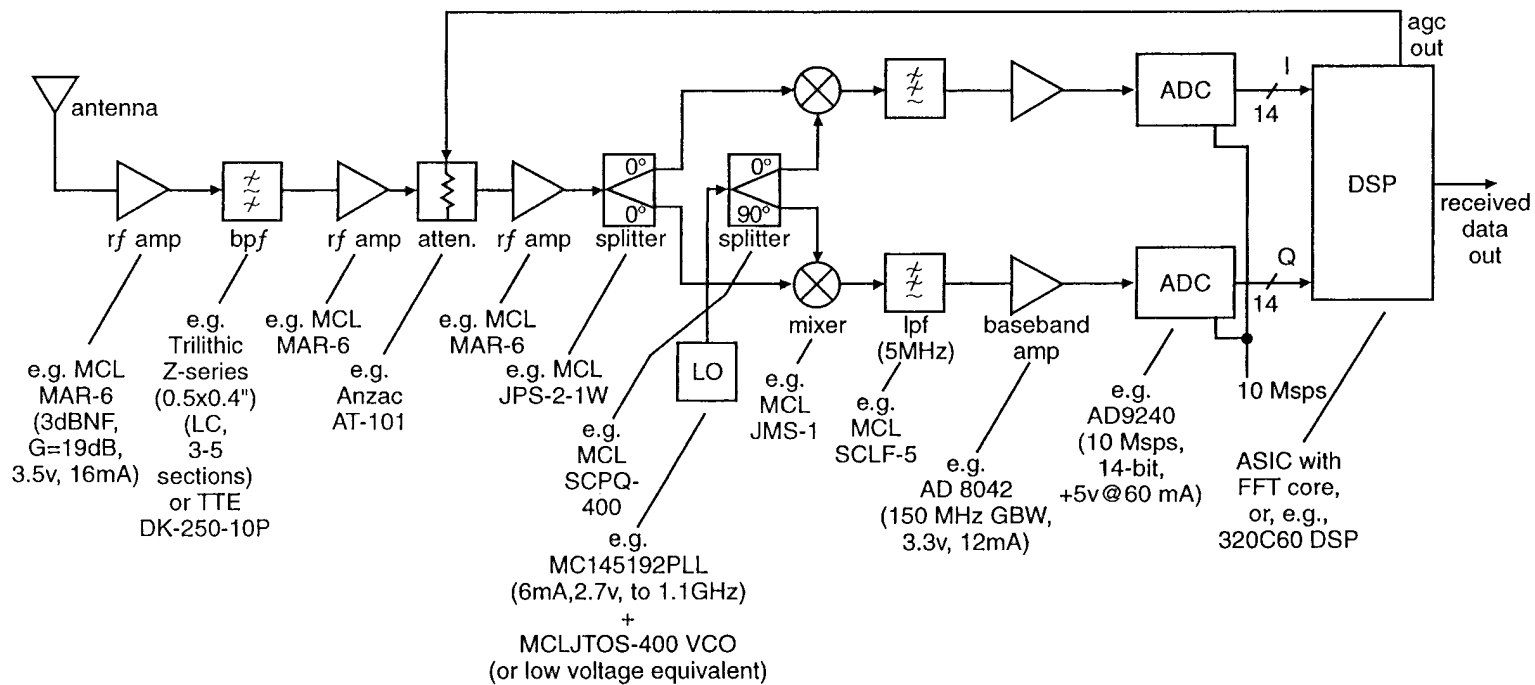


Figure 8. Hypercarrier receiver implementation using surface-mount off-the-shelf components. A 10 MHz RF bandwidth, using 1024 frequency channels of BO/OK produces a 1 Mbit/s effective data throughput after error correction.




Figure 9. Average noise power versus frequency from atmospheric, galactic, and cultural sources (Figure 1 in Chapter 29 in "Reference Data for Radio Engineers," Howard Sams, New York, 1975).

centered at f_0 . The local oscillator is conveniently implemented by PLL synthesis, for which there are numerous COTS implementations, thanks to cellular and other wireless markets.

The baseband voltages are sharply lowpass filtered at the Nyquist frequency to prevent aliasing; the quadrature voltages, roughly -30 dBm (7 mV), are amplified to ADC levels (2 V span) with fast op-amps (or, alternatively, with fixed-gain video amplifiers), then converted to 14-bit samples in a pair of low-power ADC's. Note that a 10 MHz RF bandwidth requires only 10 megasamples/sec (not 20), because they are complex. The resultant complex digital stream is Fourier transformed in a fast DSP chip: the TI 320C60 (1600 MIPS) has adequate speed, but consumes much more power (about 7 watts) than would a custom ASIC (see power, size, cost estimates below).

5.2.1 Power, Size, Cost

The power consumption of the receiver implementation of Figure 4 alone is roughly 500 mA at 3-5 V, assuming the DSP ASIC's supply current is no more than 300 mA (this is about 10% of the current used by the off-the-shelf 320C60), and that a 5 volt VCO is used for the LO. That translates to roughly an hour on a set of 4 alkaline AA cells. To this should be added the power consumption of the transmitter portion, which however is likely to operate for brief bursts only in the intended applications, thus having little impact on battery life. Battery life could be extended considerably by power-switching the power-hungry portion, namely the ADCs and DSP, upon detection of a hypercarrier signal. In that case the no-signal standby current is about 100 mA, for which 4 alkaline AA cells would provide roughly 10 hours endurance (the slow-discharge 1400 mAh capacity of alkaline AA cells is reduced to 1000 mAh at 100 mA drain, and considerably less at 500 mA drain).

The physical size of the receiver is roughly 2 inches square by $\frac{1}{2}$ inch thick, exclusive of batteries and antenna. Adding the transmitter components brings the package up to $3 \times 2 \times 0.5$ inches. The per-unit cost of components of the finished transceiver, in quantity production, is roughly \$500, assuming the ASIC unit cost not to exceed \$100.

5.2.2 Processing

The primary processing task for the receiver is to perform a 1024-point complex FFT once each $100\mu\text{s}$. This can be efficiently computed using a 5-stage radix-4 FFT algorithm. Each radix-four step involves performing a complex *twiddle-factor* multiply on each of three inputs (a total of 12 real multiplies and 6 real adds) and summing four complex values (with free 90-degree rotations) for each of the four outputs (a total of 24 real adds). To compute the $1280 = 256 \times 5$ radix four blocks, 10^4 times per second, the total computation requires 4×10^8 adds/sec and 1.5×10^8 multiplies/sec. For the signal levels in question, 16-bit integer operations with scaling shifts between the stages of the FFT are sufficient.

The transmitter needs to perform a 1024-point inverse FFT on real-valued input data. It can perform two-such real-valued FFTs using a single complex FFT by packing the two input sequences into the real and complex portions of the input and separating the even and odd components of the output sequence. Thus, the computational load for the transmitter is exactly half that of the receiver, about 3×10^8 operations/s.

The compute performance required for this task, about 6×10^8 operations/s (600Mops), is within the capabilities of an off-the-shelf signal processor, the Texas Instruments TI320C60 DSP. This chip uses a very-long-instruction-word (VLIW) organization to achieve a peak computation rate

of 1600Mops. This is sufficient computational power both to perform the required receive FFT (600Mops) and to perform a trellis decode of the convolutional code.

At considerable development expense, the power dissipation of the communication system can be reduced by replacing the off-the-shelf DSP with an ASIC designed specifically to perform the FFT processing and convolutional encoding/decoding required by this application. Such a chip would contain, for example, eight 16-bit adders, 4 16×16 multipliers, 64Kbits of static data memory, 16Kbits of microcode memory, a hardwired convolutional encoder, and hardwired trellis decoding logic. We estimate that in a modern $0.25\mu\text{m}$ CMOS process, this logic will fit on a chip measuring about 20mm^2 and will consume about 300mW.

5.2.3 Variations

One can imagine alternatives in the implementation. For example, in quantity production it may make sense to design an RF MMIC to reduce size and interconnections. One could also consider filters made with mechanical resonator or SAW technology. One could eliminate the downconversion, at modest VHF frequencies, by direct-to-baseband RF sampling; however this requires sharp RF-band filtering to prevent multiple image bands in the (undersampled) spectrum. Finally, one could use ceramic hybrid construction to reduce size.

5.3 Extra Spreading via “Tukey’s Trick”

We show here how additional spreading can be obtained *without* significantly increasing the processing burden, by using a trick of Fourier analysis

that goes back to Tukey.

On the transmitting side, to spread over a factor K greater than the baseline system, we will need to transmit a waveform that is K times more closely sampled. However, we don't want to compute the N point FFT's K times faster, so that the sampling rate r now becomes Kr . (We are assuming that N is held constant, say = 1024.) After computing each waveform by an FFT, we therefore transmit K repetitions of it, during which we compute the next waveform to be transmitted (again with K repeats.)

The effect of this scaling is to send the same number N carriers, each with the same width as before, but spread over K times more bandwidth.

On the receiving side, Tukey's trick comes into play: We must sample at the new rate of Kr , but the only operation that we do at this fast rate is a single add per sample, adding the input values onto a "wheel" of length N , until exactly KN samples are reached. In equations, we obtain a new "slow rate" series F_i that is related to the "fast rate" samples f_j by

$$F_j = \sum_{k=1}^K f_{j+kN} \quad (j = 0, \dots, N-1) \quad (9)$$

Now, the N point Fourier transform of the slow series is

$$\tilde{F}_m = \sum_{j=0}^{N-1} \exp(2\pi ijm/N) F_j \quad (10)$$

$$= \sum_{j=0}^{N-1} \sum_{k=1}^K \exp(2\pi ijm/N) f_{j+kN} \quad (11)$$

$$= \sum_{j=0}^{N-1} \sum_{k=1}^K \exp[2\pi i(j+kN)(mK)/(KN)] f_{j+kN} \quad (12)$$

$$= \sum_{n=0}^{KN-1} \exp[2\pi in(mK)/(KN)] f_n \quad (13)$$

which is seen to yield the Fourier transform values of *every K th component* of the fast series! In other words, if we only want the same number N

(and not KN) of carriers, we can do N -point FFTs at the same (slow) rate, independent of the extra spreading factor K .

Although the processing requirement for the FFTs is not increased, the sampling rate of the A/D converter *is* increased by a factor K . Therefore, to stay within the realm of current COTS components, we don't want to make K too large. A good value might be $K = 8$, which broadens the baseline system to 40 MHz bandwidth.

5.3.1 Implementation

To implement the broadened hypercarrier scheme in hardware, we need only *i*) modify the transmission protocol (broadening filters, etc., where needed), and *ii*) modify the receiver of Figure 4 by broadening the filters, increasing the sampling rate of the ADCs by the repetition factor k ($k = 8$ in the example above), and adding code in the DSP to perform pointwise addition over the k repetitions. Note that the FFTs are performed at the same rate as before. With current technology, the bottleneck is ADC speed: Analog Devices currently lists a 12-bit, 65 Msps ADC that operates from a single +3.3 V or +5 V supply (AD6640), dissipates 700 mW, and costs the same as the AD9240 (\$60 in 1k quantities). With this device one could broaden to 65 MHz, using quadrature sampling as in Figure 4; for instance, one could use a repetition factor $k = 8$, with $\tau = 125 \mu\text{s}$, $N = 1024$, $r = 8$ MHz, $\text{BW} = 64$ MHz.

This hardware implementation of the broadened hypercarrier transceiver is comparable in cost and size to the basic (unbroadened) system, though battery life is shorter owing to greater supply current for the faster ADCs; once again, power switching of those components would be advisable.

6 Reducing the Detectability of Hypercarrier

Hypercarrier's ability to operate reliably with signal levels only 10dB above background in a fading environment makes it inherently less detectable than conventional modulation methods that typically provide link margins of 30dB or more for resistance against fading. Hypercarrier also has the advantage that it spreads its power over a wide bandwidth, 10MHz for our basic system and 40MHz for the broadened system. This makes it more difficult to detect than a comparable bit-rate system that compresses its information into a more narrow frequency band.

Despite these advantages, the hypercarrier modulation scheme does require that each channel have a power level that is 10dB or so above background, making it more detectable than systems that operate with average power in each band below the noise floor. While there is no substitute for directional beams and favorable geometry, there are several ways in which the hypercarrier modulation scheme can be enhanced to reduce its detectability.

6.1 Sparse Hypercarrier

The average power per carrier can be reduced below the noise floor by coding the data so that the duty factor of each carrier is small. Using a code, for example, that maintains a maximum duty factor of 0.1 (one 1 bit every 10 bits), gives an average S/N of 0dB in each frequency band. Lower duty factors result in proportionally lower average S/N.

By encoding information in the position of the "on" bits, one can encode $\log_2 N$ bits in each "on" bit with a duty factor of $1/N$. For example, one can encode 4 bits per "on" bit in a system with a duty factor of $1/16$.

Restricting the duty factor to $1/N$ reduces the data rate by a factor $N/\log_2 N$. This can be compensated for by using substantially more frequency channels. The hardware of our broadened system can be operated for example to provide 8192 channels by eliminating the summing step and performing an 8192-point FFT. The processing requirement of this longer FFT requires approximately eight times the arithmetic performance of the baseline system, still well within the bounds of a single ASIC.

If this 8192-channel system is operated with a duty factor of $1/32$, 5 bits can be encoded in each 32-channel group each $100 \mu s$, giving a data rate of 12.8 Mbps before error-correction coding.

7 SUMMARY

We have shown that hypercarrier modulation with binary on/off keying, and with bit times longer than the whole "ringing time" of the urban environment, is potentially capable of achieving megabit per second data rates in a manner that is completely insensitive to multipath fading.

A straw-man design, using current COTS components, has been presented. While we would not want to claim that this design is exactly optimal (or even completely correct), it illustrates that brassboard tests should neither be difficult, nor expensive.

The scheme proposed here is sufficiently different from conventional high bitrate modulation methods, that it should be quite difficult for an adversary to recognize as such. Indeed, with sufficient spreading and good operational practice, it should escape detection completely.

DISTRIBUTION LIST

Director of Space and SDI Programs
SAF/AQSC
1060 Air Force Pentagon
Washington, DC 20330-1060

CMDR & Program Executive Officer
U S Army/CSSD-ZA
Strategic Defense Command
PO Box 15280
Arlington, VA 22215-0150

Superintendent
Code 1424
Attn Documents Librarian
Naval Postgraduate School
Monterey, CA 93943

Director
Technology Directorate
Office of Naval Research
Room 407
800 N. Quincy Street
Arlington, VA 20305-1000

DTIC [2]
8725 John Jay Kingman Road
Suite 0944
Fort Belvoir, VA 22060-6218

Dr Albert Brandenstein
Chief Scientist
Office of Nat'l Drug Control Policy
Executive Office of the President
Washington, DC 20500

Dr H Lee Buchanan, III
Director
DARPA/DSO
3701 North Fairfax Drive
Arlington, VA 22203-1714

Dr Collier
Chief Scientist
U S Army Strategic Defense Command
PO Box 15280
Arlington, VA 22215-0280

DARPA Library
3701 North Fairfax Drive
Arlington, VA 22209-2308

Dr Victor Demarines, Jr.
President and Chief Exec Officer
The MITRE Corporation
A210
202 Burlington Road
Bedford, MA 01730-1420

Mr Dan Flynn [5]
OSWR
Washington, DC 20505

Dr Paris Genalis
Deputy Director
OUSD(A&T)/S&TS/NW
The Pentagon, Room 3D1048
Washington, DC 20301

Dr Lawrence K. Gershwin
NIC/NIO/S&T
7E47, OHB
Washington, DC 20505

Dr Robert G Henderson
Director
JASON Program Office
The MITRE Corporation
1820 Dolley Madison Blvd
Mailstop W553
McLean, VA 22102

DISTRIBUTION LIST

Dr William E Howard III [2]
Director of Advanced Concepts &
Systems Design
The Pentagon Room 3E480
Washington, DC 20301-0103

J A S O N Library [5]
The MITRE Corporation
Mail Stop W002
1820 Dolley Madison Blvd
McLean, VA 22102

Dr Anita Jones
Department of Defense
DOD, DDR&E
The Pentagon, Room 3E1014
Washington, DC 20301

Mr. O' Dean P. Judd
Los Alamos National Laboratory
Mailstop F650
Los Alamos, NM 87545

Dr Bobby R Junker
Office of Naval Research
Code 111
800 North Quincy Street
Arlington, VA 22217

Dr Ken Kress
Office of Research and Development
809 Ames Building
Washington, DC 20505

Lt Gen, Howard W. Leaf, (Retired)
Director, Test and Evaluation
HQ USAF/TE
1650 Air Force Pentagon
Washington, DC 20330-1650

Mr. Larry Lynn
Director
DARPA/DIRO
3701 North Fairfax Drive
Arlington, VA 22203-1714

Dr. John Lyons
Director of Corporate Laboratory
US Army Laboratory Command
2800 Powder Mill Road
Adelphi, MD 20783-1145

Col Ed Mahen
DARPA/DIRO
3701 North Fairfax Drive
Arlington, VA 22203-1714

Dr. Arthur Manfredi
OSWR
Washington, DC 20505

Mr James J Mattice
Deputy Asst Secretary
(Research & Engineering)
SAF/AQ
Pentagon, Room 4D-977
Washington, DC 20330-1000

Dr George Mayer
Office of Director of Defense
Reserach and Engineering
Pentagon, Room 3D375
Washington, DC 20301-3030

Dr Bill Murphy
ORD
Washington, DC 20505

Dr Julian C Nall
Institute for Defense Analyses
1801 North Beauregard Street
Alexandria, VA 22311

DISTRIBUTION LIST

Dr Ari Patrinos
Director
Environmental Sciences Division
ER74/GTN
US Department of Energy
Washington, DC 20585

Dr Bruce Pierce
USD(A)D S
The Pentagon, Room 3D136
Washington, DC 20301-3090

Mr John Rausch [2]
Division Head 06 Department
NAVOPINTCEN
4301 Suitland Road
Washington, DC 20390

Records Resource
The MITRE Corporation
Mailstop W115
1820 Dolley Madison Blvd
McLean, VA 22102

Dr Victor H Reis
US Department of Energy
DP-1, Room 4A019
1000 Independence Ave, SW
Washington, DC 20585

Dr Fred E Saalfeld
Director
Office of Naval Research
800 North Quincy Street
Arlington, VA 22217-5000

Dr Dan Schuresko
O/DDS&T
Washington, DC 20505

Dr John Schuster
Technical Director of Submarine
and SSBN Security Program
Department of the Navy OP-02T
The Pentagon Room 4D534
Washington, DC 20350-2000

Dr Michael A Stroschio
US Army Research Office
P. O. Box 12211
Research Triangle NC27709-2211

Ambassador James Sweeney
Chief Science Advisor
USACDA
320 21st Street NW
Washington, DC 20451

Dr George W Ullrich [3]
Deputy Director
Defense Nuclear Agency
6801 Telegraph Road
Alexandria, VA 22310

Dr. David Whelan
Director
DARPA/TTO
3701 North Fairfax Drive
Arlington, VA 22203-1714

Dr Edward C Whitman
Dep Assistant Secretary of the Navy
C3I Electronic Warfare & Space
Department of the Navy
The Pentagon 4D745
Washington, DC 20350-5000