

A Computationally Efficient Adaptive Transceiver for High-Speed Digital Subscriber Lines

J. S. Chow J. C. Tu* J. M. Cioffi†

Information Systems Laboratory, Stanford, CA 94305

Abstract

This paper presents the feasibility of the implementation of a Multichannel IDFT/DFT transceiver for crosstalk limited Carrier Serving Area (CSA) subscriber loops. This transceiver can provide reliable T1 transmission in the presence of crosstalk noise on most CSA loops. We examine the complexity of the transceiver design in terms of speed or computing power. We also note several advantages of building this system in practice: channel response is estimated during call initialization, simple transmitter, short-length equalizer, and absence of error propagation. We conclude that, without a Viterbi decoder or an echo canceller, we can implement the complete receiver with under 40 MIPS of computing power, using a fairly high signaling rate of 640 kHz, on conventional DSP chips. A similar performing Decision Feedback QAM system (or Tomlinson Precoded system) is found to require around 100 MIPS, even after significant complexity reduction. The Viterbi detector, however, requires custom design with either multichannel or QAM system, which is not discussed here.

1 Introduction

In [1] we present a computationally efficient multichannel transceiver design which is based on the Inverse Discrete Fourier Transform (IDFT) and the Discrete Fourier Transform (DFT) vectors, and is the extension of the work of [2]. The performance of this system is analyzed in [1], and is found to be close to the performance of several optimal systems [3] - [6].

In this paper, we look at the implementation issues of this design. We first describe the carrier serving area (CSA) loops in Section 2. Then in Section 3 we focus on the complexity of the actual implementation. Due to the non-minimum phase nature of some of the channels that we studied, a variation of the original transceiver design of [1] is used, and this transceiver works regardless where the channel zeros are located. Complexity in terms of Million Instructions Per Second (MIPS) are presented in detail for

each block of our transceiver design, and an efficient adapting algorithm is presented to reduce the computations required.

We take advantage of the well-known Fast Fourier Transform (FFT) algorithm, which results in the system complexity being proportional to $N \log_2 N$, where N is the block-length of the IDFT/DFT vectors used for multichannel transmission. We are able to separate two tasks: optimizing throughput in the presence of crosstalk noise, and converting the channel to partial response, by using multichannel approach. Thus we can use a relatively short equalizer to reduce the computations required compared to a scalar Decision Feedback Equalizer (DFE), as discussed in Section 4.

In addition, we compensate for slow variation of the channel response adaptively, and we do not suffer from the error propagation of a DFE system that can significantly reduce performance. We conclude that this IDFT/DFT multichannel transceiver, without considering coding and decoding, can be implemented and operated economically with current technology. However, the cost of the Viterbi decoder in the receiver is significant, and methods for reducing this complexity is under current investigation [7].

2 CSA Loop Parameters

The CSA is a subset of local loops between the telephone company end office and the subscriber. The CSA presently constitutes 75-80% of all twisted-pair subscriber loops and excludes loops with excessive lengths and small wire gauges [3]. This subset of local loops is characterized by very long post-cursor intersymbol interference (ISI). The dominant noise impairment is crosstalk noise from other twisted-pair loops within the bundle.

A set of hypothetical local loops are provided by Bell Communications Research¹. These loops are described briefly in [1].

With a sampling rate of 640 kHz, we obtained pole-zero models of the pulse responses of different channels using system identification tools. We showed in [1] that all the pole-zero models closely match the channel pulse responses, and thus use of the models for further study is justified. We found that with about eight zeros and ten poles, we have

¹We would like to thank Pedro Crespo and Russell Hsing of Bell Communications Research for the data on CSA subscriber loop channels

*Supported by Hughes Aircraft Company Doctoral Fellowship

†Supported in part by NSF grant number MIP 86-57266, by a grant from Bell Communication Research and by an internal seed grant from Stanford University's Center for Integrated Systems

348.5.1.

1750

CH2829-0/90/0000-1750 \$1.00 © 1990 IEEE

a good model for each of the channels at a sampling rate of 640 kHz. In addition, we also noted that most channels are non-minimum phase, i.e., at least some of the zeros are outside the unit circle.

3 Complexity of IDFT/DFT Signaling

The block diagram of the IDFT/DFT Multichannel transceiver is shown in Figure 1. This structure is based on the work of [2], which uses cyclic extension to construct ISI-free parallel subchannels using vectors of IDFT/DFT, except that we implement an adaptive equalizer at the receiver so as to convert the IIR channel to a partial response (FIR) channel. We could avoid performing cyclic extension by equalizing the channel response to a flat (Nyquist) shape. However, as mentioned in Section 2, some of the channels are non-minimum phase. Inverting the zeros using an IIR filter results in unstable system. An alternative would be to use an FIR filter to cancel the zeros, but the length of the filter and the delay involved become excessive for high rate data transmission. As a result, we accept the rate loss by using cyclic extension and optimize the throughput on the resulting partial response channel.




Figure 1: Block Diagram of the IDFT/DFT Transceiver

During the initial startup, a pole-zero estimation (such as an ARMA estimation) of the channel is performed. We then set the number of cyclically extended samples to be ν , the number of estimated channel zeros, and set the initial tap values of the adaptive equalizer to be the same as the estimated channel pole polynomial. The signal-to-noise-ratio (SNR) is obtained for each frequency bin and sent back to the transmitter to allocate the power and bits strategically. The performance of this transceiver is evaluated using computer simulations, and is found to be close to the performance of the slightly different multichannel transceiver in [1].

Several assumptions are made in computing the total complexity:

- Complexity is measured in terms of Million Instructions Per Second (MIPS).
- Signaling rate is equal to 640 k symbol/sec.
- A real add, a real multiply, or a real multiply followed by a real add (multiply-accumulate) is counted as one

instruction; a complex multiply/divide is counted as four instructions.

- The number of instructions is assumed to be $N \log_2 N$ for an N -point real input FFT (or equivalently N -point real output IFFT), where N is less than or equal to 128, otherwise, the number of instructions is equal to $1.5N \log_2 N$. (N is a power of 2.)²
- Computations required in timing recovery and decision block, which are not significant, are not included.
- The echo canceller, whose complexity (or even existence) depends largely on whether full duplex or half duplex is desired, is not addressed here.
- Complexity results in coding/decoding is not included, but is significant if a trellis code is used.

3.1 Transmitter Complexity




Figure 2: Block Diagram of the Transmitter

One of the main advantages of the IDFT/DFT Multichannel transceiver is the simplicity of the transmitter, as shown in Figure 2. Regardless of the channel response, fixed IDFT vectors are used to modulate the input signal, and this modulation is implemented using an N -point IFFT. The resulting data is then cyclically extended by ν . Without considering coding, the complexity of the transmitter is shown in Table 1, where bit mapping from input bits to modulator input is negligible in comparison, and thus is not included.

Note that because of the cyclic extension, we suffer a rate loss of $\nu/(N + \nu)$ for a fixed sampling rate of 640 kHz. However, if we increase the signaling rate so that we have the same amount of data transmitted in a fixed period of time as is transmitted without any cyclic extension, then the numbers given in Table 1 (and also in Table 2 and Table 3) are the complexity counts including the effects of ν . These numbers are invariant to different values of ν . In practice, signaling rate is rarely changed, so instead we take the penalty of a rate loss of $\nu/(N + \nu)$. A typical value of ν , as described in Section 2, is approximately between 5 and 12, and from [1] we know that a good blocklength N would be in the order of 512. This translates to a rate loss of at most about 2.3%, and thus is negligible. In return, slightly lower complexity results because the data are processed in a time frame of $N + \nu$ samples, instead of N .

²Using Motorola's DSP96002, the instruction count is about $2.4N \log_2 N$, and this is the best currently available, but we presuppose yet better efficiency in implementing the FFT with $1.5N \log_2 N$.

phase one, and thus the number of feedforward taps and the number of feedback taps are much larger in order to get good performance. Results in [3] show that approximately 128 taps are needed for the CSA loops we studied. The transmitter, without encoding complexity, would have low complexity. However, the scalar DFE on the receiver requires about 165 MIPS of computing power. Even if we do not update the weight every time, the complexity of the receiver, without considering decoding, would still be close to 100 MIPS. Another practical drawback of using scalar DFE is the problem of error propagation of the feedback taps. Error propagation degrades the achievable throughput of the system, although an equivalent Tomlinson Precoder can be used to address this problem; complexity remains the same. We should also note that the use of FFT's in DFE systems does not appear to be attractive [10].

One drawback of our system is the extra storage required. However, in high rate data transmission system, the most important constraint appears to be the speed, not the storage. Thus by sacrificing extra memory, we should be able to have a practical, implementable real time system.

5 Conclusions

We have presented the basic structure and the complexity of implementing the IDFT/DFT multichannel transceiver. We describe an efficient adaptive equalizer which serves to convert the channel response to an FIR one, and by using multitone approach we allocate different energy levels into different frequency bins to optimize throughput.

The transmitter is very simple to implement since the modulating vectors are fixed. Without considering decoding, we can implement the complete receiver with less than 40 MIPS of computing power, which is well within the range of current technology.

With some custom design chips and exploiting the possibility of parallel processing, we should be able to build a complete, real-time transceiver system with only a few VLSI chips. One bottleneck in the receiver is the Viterbi decoder, which may take around 1000 MIPS using a standard DSP chip. However, due to the special structure of the Viterbi algorithm, it may be possible to reduce the complexity of the decoder by a significant amount by switching to a custom VLSI chip [7].

6 Acknowledgements

The authors would like to thank the following people in alphabetical order for helpful discussions and suggestions throughout the development of this paper: J. Aslanis, P. Chow, G. Dudevoir, M. Ho, S. Kasturia, H. Lou, T. Ruiz, and R. Ziegler of Stanford University, P. Crespo, R. Hsing, and J. Lechleider of Bell Communication Research, J. Bingham of Telebit, and N. Jablon of AT&T.

References

- [1] J. Tu, J. Chow, G. Dudevoir, and J. Cioffi, "Crosstalk-limited performance of a computationally efficient multichannel transceiver for high rate digital subscriber lines," in *1989 Global Telecommunications Conference*, (Dallas), November 1989.
- [2] A. Ruiz, J. Cioffi, and S. Kasturia, "Frequency-specified coset codes for spectrally-constrained channels," March 1988. submitted, *IEEE Transactions on Communications*.
- [3] G. Dudevoir, J. Chow, S. Kasturia, and J. Cioffi, "Combined equalization and coding for T1 data rates on carrier serving area subscriber loops," in *1989 International Conference on Communications*, (Boston), June 1989.
- [4] D. Falconer and V. Joshi, "Channel capacity bounds for the subscriber loop," in *IEEE Pacific Rim Conference on Communications, Computers and Signal Processing*, June 1987.
- [5] I. Kalet and N. Zervos, "Multitone and single-tone DFE — The connection,". Submitted to *IEEE Trans. on Communications*.
- [6] J. Aslanis and J. Cioffi, "Capacity and cutoff rate of the digital subscriber loop with near end crosstalk noise," in *1989 International Conference on Communications*, (Boston, MA), June 1989.
- [7] H. Lou, "A specialized processor architecture for viterbi decoding," in *1990 IEEE High-Speed Digital Subscriber Lines Workshop*, (Parsippany, New Jersey), March 1990.
- [8] G. Clark, S. Mitra, and S. Parker, "Block implementation of adaptive digital filters," *IEEE trans on ASSP*, vol. 29, pp. 744-752, June 1981.
- [9] S. Kasturia, J. Aslanis, and J. Cioffi, "Vector coding for partial-response channels," February 1988. submitted to *IEEE Trans. on Information Theory*.
- [10] N. Jablon, "Complexity of frequency-domain adaptive filtering of data modems," in *1989 Asilomar Conference on Signals, Systems and Computers*, (Monterey), November 1989.

Blocklength (N)	64	128	256	512	1024
Complexity ^a (in MIPS)	3.8 (9.2)	4.5 (10.8)	7.7 (12.3)	8.6 (13.8)	9.6 (15.4)

(Note: ν is assumed to be zero.)

^aThe number in parenthesis assumes using Motorola's DSP96002

Table 1: IFFT Transmitter Complexity at 640 kHz sampling rate

Complexity ^a (in MIPS)	L=5	L=10	L=16	L=32	L=64
N=64	19.9 (30.6)	26.3 (37.1)	34.1 (44.8)	54.7 (65.5)	96.0 (106.8)
N=128	21.1 (33.7)	27.6 (40.1)	35.3 (47.8)	55.8 (68.4)	97.0 (109.5)
N=256	27.5 (36.7)	33.9 (43.2)	41.6 (50.9)	62.2 (71.4)	103.2 (112.4)
N=512	29.4 (39.8)	35.9 (46.2)	43.5 (53.9)	64.0 (74.4)	105.0 (115.4)
N=1024	31.4 (42.9)	37.8 (49.3)	45.5 (57.0)	65.9 (77.5)	106.9 (118.4)

(Note: ν is assumed to be zero.)

^aThe number in parenthesis assumes using Motorola's DSP96002

Table 2: Receiver Complexity with Full Block LMS

Complexity ^a (in MIPS)	L=5	L=10	L=16	L=32	L=64
N=64	16.4 (27.1)	20.4 (31.2)	25.9 (36.7)	44.2 (54.9)	96.0 (106.8)
N=128	17.5 (30.0)	21.1 (33.6)	25.8 (38.3)	40.0 (52.5)	76.2 (88.7)
N=256	23.8 (33.0)	27.2 (36.4)	31.4 (40.7)	43.7 (52.9)	72.0 (81.2)
N=512	25.6 (36.0)	29.0 (39.3)	33.0 (43.4)	44.2 (54.6)	68.6 (79.0)
N=1024	27.5 (39.1)	30.8 (42.3)	34.7 (46.3)	45.5 (57.0)	67.9 (79.4)

(Note: ν is assumed to be zero.)

^aThe number in parenthesis assumes using Motorola's DSP96002

Table 3: Receiver Complexity with Block LMS with efficient updating (M=L)

Complexity ^a (in MIPS)	L=5	L=10	L=16	L=32	L=64
$\nu=5$	25.4 (35.7)	28.7 (38.9)	32.7 (42.9)	43.8 (54.1)	68.0 (78.2)
$\nu=8$	25.2 (35.5)	28.5 (38.7)	32.5 (42.7)	43.6 (53.8)	67.6 (77.8)
$\nu=12$	25.1 (35.2)	28.3 (38.4)	32.2 (42.4)	43.2 (53.4)	67.1 (77.2)
$\nu=16$	24.9 (34.9)	28.1 (38.1)	32.0 (42.1)	42.9 (53.0)	66.6 (76.6)

^aThe number in parenthesis assumes using Motorola's DSP96002

Table 4: Receiver Complexity with Block LMS with efficient updating (N=512, M=L)

348.5.5.