

Demonstration of Uncoordinated Multiple Access in Optical Communications

H. Chan¹, A. Vila Casado¹, J. Basak¹, M. Griot¹, W. Weng¹, R. Wesel¹, B. Jalali¹, E. Yablonovitch¹, I. Verbauwhede^{1,2}
{herwin, avila, juthika, mgriot, wenyen, wesel, jalali, eliy, ingrid}@ee.ucla.edu

¹Electrical Engineering Dept.
UC Los Angeles, USA

²Dept. ESAT/SCD-COSIC
K.U.Leuven, Belgium

ABSTRACT

Optical communications has been demonstrated to offer very high bandwidths. However, complexities in the use and setup of an uncoordinated multiple access network curtail its use in local area networks and is deemed to be less useful compared to the Ethernet. In this work, we demonstrate a system with a convenient scheme for uncoordinated multiple access, maintained at the high optical transmission rates. The system uses a Non-linear trellis code together with a Reed Solomon code at optical throughputs. A six user demonstration was implemented with each node implemented on a Xilinx VirtexII-Pro FPGA. Commercial off the shelf components were used to build both the electrical and optical parts of the system. We successfully demonstrate an optical network with a bandwidth of over 7 Gbps, which is an order of magnitude above that of current Ethernet technology.

Categories and Subject Descriptors

B.6 [Hardware]: Logic Design; C.2.5 [Computer System Organization]: Computer Communication Networks, Local and Wide Area Networks.

General Terms

Algorithms, Performance, Design..

Keywords

Optical Network, Nonlinear Trellis Code, Multiple Access Network, Uncoordinated Access Network.

1. INTRODUCTION

For a particular frequency, optical channels promise data rates of up to hundreds of gigabits per second. This is because there is very little noise or signal degradation in the transmission of light through modern fiber-optic cables. These properties of optics have been well utilized to form the backbone of global networks such as the internet and telephone networks. Communications among the optical transmission stations approach 100% efficiency

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

DAC'06, Month 1-2, 2004, San Francisco, CA, USA.

Copyright 2006 ACM 1-58113-000-0/00/0004...\$5.00.

(requiring no redundancy for error correction) by means of WDMA and TDMA techniques. This efficiency, however, requires considerable coordination between all the communicating nodes.

In the local area network domain, optical networks have had limited success. Though the optical token ring (FDDI) network promises higher bandwidth [1], Ethernet networks are significantly more popular. This success is due to the ease in which a network can be set up. Nodes on the network can be added and removed from the network dynamically without disruption of network communications. In addition, there is no coordination between the nodes when a certain node wants to transmit data; the node merely starts to transmit and protocols determines what to do when a collision occurs.

Though there are efforts to implement Ethernet on optical networks, aggregate throughput performance is fundamentally limited by collision of data [2]. We demonstrate the desirable properties of Ethernet by marrying the high-bandwidth properties of optical networks with the flexibility of Ethernet. Collisions are avoided by careful design of channel codes. The uncoordinated multiple access properties will be provided by a set of novel channel codes, which guarantee that data can be decoded at optical bit error rates i.e. BER<10e-9. This bit error rate performance will be maintained even in the presence of other (interference) transmissions. Even with uncoordinated access, we are able to achieve 30% efficiency by treating the interference users as noise.

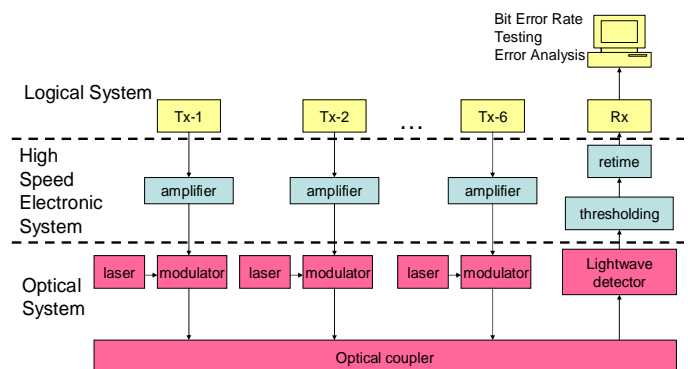


Figure 1. High-level view of demonstration system

This paper describes the design and implementation of a system that demonstrates an uncoordinated multiple access optical network. The system features six nodes simultaneously transmitting into an optical channel. A receiver node takes the aggregate signal and decodes a single user. Figure 1 shows the system block diagram of our system. Each transmitting node is comprised of an FPGA, which codes the data, a laser, which provides the carrier for the coded data, and a modulator, which combines the two together and puts the data on the channel.

The goal of our demonstration system is to show correct operation of six users on a single optically coupled network. The channel bit rate of 2 Gbps is divided among the 6 nodes using novel channel coding techniques. This will give an uncoded useful data bandwidth of 93 Mbps for each node.

Though the system was built with commercial off the shelf components, the design involves several domains. Section 2 describes the design process of the channel code. Section 3 describes how these algorithms were implemented onto a FPGA platform. Section 4 describes the optical system and how the logical system interfaces with it. Section 5 analyzes the results of the final demonstration system setup and Section 6 concludes the paper.

2. CODE DESIGN

A simple communications model that can describe the multiple-user optical channel with non-coherent combining is the OR channel. In this channel, if all users transmit a zero, then the channel output is a zero. However, if even one user transmits a one, then the channel output is a one. Information theory tells us that the maximum sum-rate (the sum of the rates of all the transmitters in the system) of the OR channel is 1 information bit per received data bit. However, for a simple uncoordinated access decoder, other users must be treated as noise. This transforms the OR channel into the Z-Channel, and lowers the symmetric sum rate to less than one. The remaining tasks consists of selecting the multiple access strategy and the channel code that will be used.

For uncoordinated multiple-access, Interleaver-Division Multiple-Access (IDMA) [3, 4] is a promising approach. With IDMA, every user has the same channel code, but each user's code bits are permuted using a unique randomly drawn interleaver. The receiver is assumed to know the interleaver of the desired user. With IDMA in the OR multiple access channel (MAC), a receiver should see the desired signal corrupted by a memoryless Z-channel. We performed simulations comparing a Non-Linear Trellis Coded Modulation (NL-TCM) code under two channels: 1) a 6 user OR-MAC channel using IDMA and 2) the equivalent Z-channel that the receiver would see if the errors were not generated by codewords but by random errors. The performance was the same. Thus, in the context of IDMA, the remaining challenge is the design of a good code with the desired ones density.

In order to achieve the maximum symmetric sum-rate where each user sees a Z-channel, the channel code must have a particular average density of ones p_1 . The channel code used is a Non-Linear Trellis Coded Modulation (NL-TCM) code. This novel code provides us with the appropriate information rate and density of ones. A Viterbi decoder allows a simple and fast

decoding of NL-TCM. A brief description of the design of these codes is presented in the following subsections.

2.1 Directional Hamming Distance

Regular convolutional codes are design so that the Hamming distance between codewords is maximized. Hamming distance is the number of bits that differ between the codewords. This distance is directly associated with the number of errors such a code can decode. In the Z-Channel, a transmitted 1 will always induce a received 1. Thus, to make a decoding error, the decoder must see ones in all the bit positions where the incorrect codeword has ones. This implies that a new definition of distance is required. Let us define the **directional Hamming distance** $d_D(c_1, c_2)$ the number of ones that we have to add to c_1 so that all ones of codeword c_2 are ones in the received word.

Given that the purpose of the design is to maximize this directional distance, the safest definition of distance between branches would be

$$D_{i,j} = \min[d_D(c_i, c_j), d_D(c_j, c_i)],$$

which is the branch-wise metric that will be maximized in our design. By taking the minimum between the two directional distances as the metric to maximize, we seek to maximize the minimum directional distance, albeit in a greedy fashion.

With this branch-wise metric, codewords with equal Hamming weights produce larger values than codewords with different Hamming weights, so we will assign output values to the trellis branches with as similar Hamming weight as possible, preferably equal.

2.2 NL-TCM Code Design

The trellis code design consists on assigning output values to the branches of the trellis code. Those outputs have to maintain the desired average density of ones p_1 . Our goal is to maximize the minimum directional distance d_{min} using the greedy pairwise metric.




Figure 2. Basic sub-graph of the trellis diagram

The idea is to apply Ungerboeck's rules [5] in the context of our pairwise metric. To explain how to do that, consider the subgraph shown in Figure 2. That figure shows the subgraph that composes every trellis diagram that results from a conventional feed-forward convolutional encoder with one input. The branches produced by an input bit equal to 0 for both states go to the same state, and the same happens with an input bit of 1. Thus, we have in the same subgraph all the splits and merges produced by those four branches. For a 2^v -state encoder, there are 2^{v-1} subgraphs of this type.

Ungerboeck's rules consists maximize the distance between branches splitting from a state (splits) and branches merging to a

same state (merges). Ungerboeck's rules can be extended more deeply into the trellis. For instance, if we could choose all branches to have distance of at least 1 between them, the distance of different paths would be increased by at least 1 in each step. Moreover, we can try to maximize the distance not only of the splits, but also the distance between the four branches coming from a split in the previous step (see Figure 3), and the eight branches that come from a split two steps before, and so on. And we can move backwards from a merge trying to maximize the distance between the previous four branches, and so on (see Figure 3).




Figure 3. a) Four paths that start on the same state in two trellis sections. b) Four paths that arrive to the same state in two trellis sections. Branches are labeled with the input bits that induce traversal of the branch

Therefore, the design consists on assigning output values that satisfy the target density p_1 of ones and satisfy the extended Ungerboeck's rules. Further details about the design of the codes can be found in [6].

2.3 Block Code

Optical systems typically deliver a very low BER. In order to maintain this BER, the rate of the NL-TCM channel code would have to be very low. A better solution is found taking into account the distribution of the erred bits in a transmitted stream after the NL-TCM decoding. It is well known that in Viterbi decoding, a small number of erred bits cluster together. Thus, a high rate block code that can correct few symbol errors can be attached as an outer code, dramatically lowering the BER.

A concatenation of the rate-1/20 NL-TCM code with a (255 bytes, 237 bytes) Reed-Solomon code has been tested for the 6-user OR-MAC scenario. The rate of this code is $(237/255) \times (1/20) = 0.0465$. The simulated BER is $2.5e-10$. For six users, the sum-rate is $6 \times 0.0465 = 0.279$.

3. CODE IMPLEMENTATION

The coding algorithms were implemented on the Xilinx VirtexII-Pro FPGA [7]. The implementation dataflow block diagram is shown in Figure 5. Data to be transmitted is first encoded with a reed solomon block code. The output bits of this block are encoded with the trellis code and then passed on to the interleaver. Finally, the resulting bits are sent to the high speed serial transceiver (Rocket I/O) to be sent off-chip. Figure 4 shows that the trellis code together with IDMA provide the uncoordinated access properties of our system and are able to bring the bit error rate (BER) to about $1e-5$.




Figure 4. Bit error rate of NL-TCM codes versus the number of users

All the nodes of the system use the same code, but the IDMA interleavers are used to ensure that the coded bit patterns do not look the same in the optical channel. Finally, the Reed-Solomon block code is used to further reduce the BER to less than $1e-9$. In addition to these blocks, synchronization blocks ensure that the received bits are aligned properly so that decoding can be performed correctly.

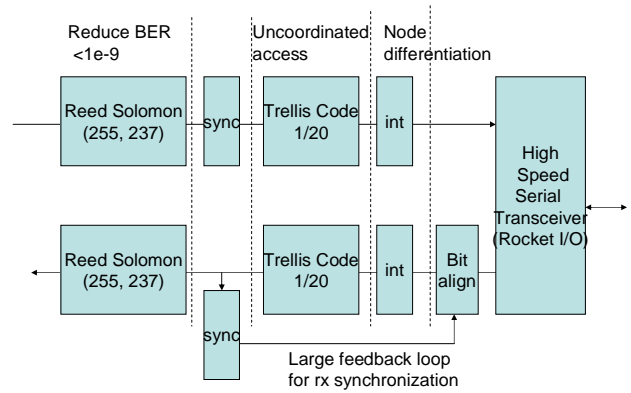


Figure 5. Coding implementation dataflow

3.1 Interleavers

The output of an interleaver is a reordering of the input bits. In our system, since a transmitter is defined by the interleaver pattern, the most desirable pattern is one that is generated randomly. This implementation option is not feasible in our demonstration system because it will be too large and too slow. For the desired channel rate of 2 Gbps, using the 20 bit interface of the Rocket I/O, our target operating frequency for the interleaver is 100MHz. A critical path of 7.7ns easily achieves this requirement.

For our 1600 bit interleaver implementation, we adopted a randomized write-by-row, read-by-column scheme. As seen in Figure 6, data can be broken into square blocks of 400 bits. Each of the 20 rows and columns are indexed. Groups of 20 incoming bits are written to a randomly indexed row. When the data block is filled, the bits are read out of the block one column at a time in a random order.

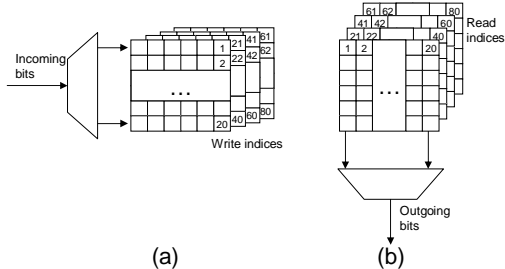


Figure 6. Indexed write-by-row, read-by-column interleaver

This scheme provides us with enough randomness to operate on the optical channel. In our interleaver design, 4 square blocks of 400 bits are used, giving us a total of 80 indexed locations. This corresponds to a design space of $(80!)^2 > 1e+237$ possible interleavers to choose from.

3.2 NL-TCM Code

The NL-TCM rate 1/20 code from the previous section is chosen to be implemented. This is a convenient choice since the interface with the high speed transceiver is 20 bits also. Figure 7 shows the Viterbi decoder architecture used in the receiver.

Due to the high throughput requirements, the decoder has a fully parallel 64 state implementation. For each incoming 20 bit codeword, 128 individual path metrics are computed simultaneously then compared and accumulated. To meet timing requirements, the design is heavily pipelined. The traceback depth of our implementation is 35.

The calculated branch metric of our Viterbi decoder is different from traditional implementations. The branch metric is calculated by comparing the received codeword with the expected codeword. Traditionally, the metric is simply the number of bits that do not match. In our case, since the OR channel does not allow a '1' to turn into a '0', when such a case is encountered, the branch metric will be automatically set to a value of 20.

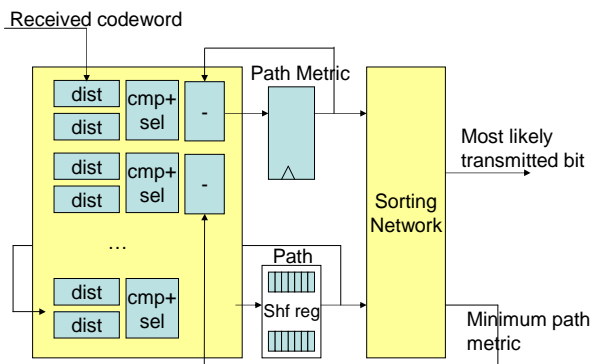


Figure 7. Architecture of Viterbi parallel decoder

After all the branch metrics are computed, a sorting network is used to select the path with the smallest accumulated metric. A minimum time sorting network based on Batcher's odd-even merging algorithm [8] is used. The minimum path metric is also fed back to the Viterbi decoder and subtracted from the accumulated path metric. This is to ensure that the register

values do not overflow.

For 20 bit codewords, a critical path of 10ns is desired to achieve the 2 Gbps channel rate. However, a critical path of 10.3ns is reported after place and route. This path is found in the feedback path of the path metric calculations and therefore cannot be further pipelined.

3.3 Reed Solomon Code

When the Trellis decoder block makes an error, the errors usually come in a burst of a few bits at a time. A Reed Solomon (RS) code is a block code that operates on bytes at a time. This makes it a very good choice to correct the residual errors and bring the final BER to below $1e-9$. A standard (255,237) RS code was selected.

Since timing is not critical in this block, a standard architecture design is used [9]. The syndromes of the input data block are first calculated. The results are then used to calculate the error locator polynomial using Berlekamp's algorithm. The Chien algorithm is used to find the roots of the error locator polynomial and these roots provide the location of the errors. Finally, the magnitudes of the errors are captured.

The data rate at the output of the NL-TCM decoder is 100Mbps. Since the Reed Solomon code operates on data blocks of 255 bytes (2040 bits), the time budget for the RS decoder is 20.4us. We clocked the module at 50MHz, and at the worse case the decoding operation takes 856 cycles (17.1us) to complete.

3.4 Implementation Results

The system blocks were implemented on the VirtexII-Pro FPGA from Xilinx. Table 1 summarizes the size various blocks in the design. The critical period is given for the transmitter and receiver.

Table 1. Size and speed of transmitter and receiver blocks

	Area (slices)	Critical period (ns)
Transmitter		
Reed Solomon encode	189	
Synchronization	32	
NL-TCM encode	34	
Interleaver	3387	7.7
Receiver		
Reed Solomon decode	3686	
Synchronization / bit alignment	162	
NL-TCM decode (Viterbi)	10504	10.3
Interleaver	3387	

The transmitter is implemented on the VirtexII Pro XC2VP20 FPGA which contains 9,230 slices of logic. Each transmitter design occupies 40% of the available area. The receiver is a significantly larger design and is implemented on the XC2VP50 which has a capacity of 23,616 slices. The receiver design occupies 70% of the available area.

4. ELECTRICAL / OPTICAL SYSTEM

The design of the nonlinear trellis codes is critically dependent on the assumption of incoherent addition of data coming through the

individual users. In the optical domain, this implies that there should be no destructive interference between light coming from the different users. If this condition holds true, a combination of '1's from any two users would never result in a '0' and would constructively add to a '1'. Using different lasers for the different channels ensures the required coherence between each channel and eliminates the chance of destructive interference.

Continuous wave (CW) lasers with 1550 nm optical wavelength are used in this system. The electrical-to-optical (E/O) conversion is achieved using a photodetector, which is in fact an intensity (not amplitude) detector, and serves as a mixer for the six optical signals.

Since they are oscillators, lasers contain some phase noise around its carrier frequency [10]. Because of this property, there is a chance that 2 lasers being within ~0.03 nm (4 GHz when centered around 1550 nm), mixing at the photodetector would result in the phase noise associated with the lasers to be brought down to the baseband. This would cause the system to be affected by a large amount of low frequency phase noise, which could make operation of the system impossible. However, for the ITU wavelength grid of 50 GHz, 6 lasers separated by 4 GHz each keep us well within the available bandwidth.

The optical multiple access channels designed here are based on intensity-modulated direct-detection (IM-DD) links [11]. External intensity modulation of the light is performed at the transmitter end using LiNbO₃ Mach-Zehnder (MZM) modulators. These modulators are widely used in research on radio-over-fiber links [12]. Direct detection implies intensity detection by a photodetector at the receiver end, with no information retained about the phase of the light that was transmitted through the channel. This section describes the detailed design of the transmitter and receiver end required for the system. Figure 8 is a block diagram of the system being described below.

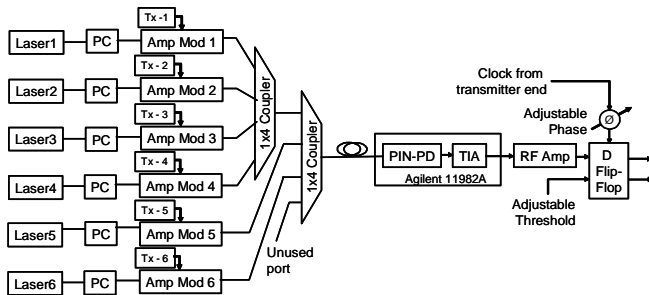


Figure 8. Electrical / optical system architecture

4.1 Optical Transmitter

The exact nature of the lasers to be used in this experiment was not critical and we used different kinds of lasers depending on the availability. Four lasers derived from the four channels of a Santec laser (ECL-200) and two other lasers from JDSU (CQF938/500) and Fujitsu (FLD5F7CZ-J) were used as the six optical transmitters. Six MZM modulators, four from Sumitomo (T.DE1.5-10-P) and two from EOspace (10 Gbps 1x2 Dual-output modulator) were used for the intensity modulation. Since these modulators are dependent on the polarization of the optical input, polarization controllers (PC) were placed in the optical path before being input into them. The polarization was adjusted to give

maximum optical output from the MZM modulators.

4.2 Optical Receiver

Two 1x4 optical couplers were used to combine the 6 optical channels. Direct detection of the transmitted data is performed by an HP 11982A lightwave converter which consist of a p-i-n photodetector (PIN-PD) followed by a Transimpedance Amplifier (TIA). The detected RF signal is a result of the data from all the 6 users added together. Since the HP 11982A has no limiting characteristics, the amplitude of the output is proportional to the number of users transmitting a "1". However, this signal needs to be converted into a binary form for it to be recognizable to the FPGA receiver.

A D flip-flop following the lightwave converter is used to convert this multilevel signal into a binary signal. The D flip-flop samples the input data at every positive edge of the clock fed into it. As mentioned earlier, the transmitters are all asynchronous to each other. Thus the receiver needs to have the capability to be synchronized to any desired user. For this purpose, the clock input in to the D-flip flop is followed by an adjustable RF phase delay line, as shown in Figure 8. This changes the relative phase between the clock and the signal. Thus the received signal is retimed in synchronization with any given user.

As the data is sampled by the D flip-flop, it is compared to an adjustable threshold. Depending on its value relative to the threshold, the multi-level photodetected output is converted into a binary signal. Thus, the D flip-flop also performs the function of regeneration. The output binary signal is designed to have 'low' and 'high' voltage levels that are recognizable by the FPGA receiver.

5. RESULTS

Pictures of the demonstration setup is shown in Figure 9. Figure 9(a) shows the laser sources. Figure 9(b) shows how two of the FPGA transmitters are connected to the optical network. The light travels from left to right and passes through the polarization controllers (in black) and to the optical modulators (silver boxes on the right). Figure 9(c) shows the computer used to display the bit error rate and the oscilloscope to look at the raw received waveform.

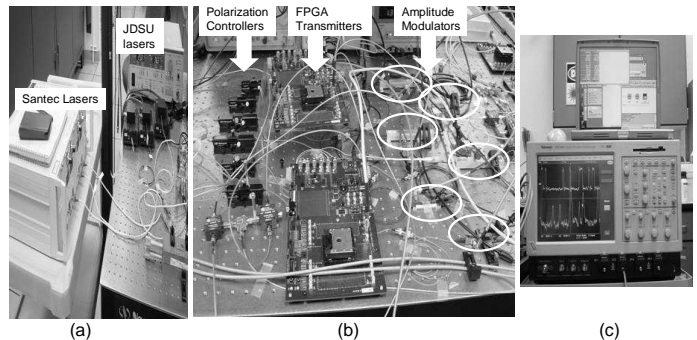


Figure 9. Demonstration setup components: a) laser sources b) FPGA's and laser modulation b) system monitor and measurement

Figure 10(a) shows the raw received waveform for the case of four

simultaneous transmitters. After the thresholder and D flip flop circuit, the result is shown in Figure 10(b). This is the waveform that is given to the receiver FPGA to decode.




Figure 10 Four user case of receive signal threshold and retiming

Testing of the system proceeded in the following manner. The desired user transmits a constant pattern in which the receiver FPGA is able to detect. This channel is activated first, and the threshold and sampling moment is adjusted to the correct point. Each of the other FPGA transmitters are set to transmit random coded data. This interference is added to the optical channel one at a time so that the threshold and sampling time may be tweaked. This proceeds until all six transmitters are simultaneously transmitting on the optical channel. The results of this are presented in Table 2.

Table 2. System results

Channel rate	1.2 Gbps * 6 = 7.2 Gbps		
Data rate	60 Mbps * 6 = 3.6 Gbps		
Users	Bits tested	Errors found	Measured BER
1	1.5e11	0	< 6.4e-12
2	4.6e10	0	< 2.2e-11
3	1.2e9	0	< 8.3e-10

Interfacing problems with the Rocket I/O transceiver we were only able to demonstrate a system with 3 users at a channel bit rate of 1.2 Gbps. This performance degradation can be attributed to two main factors: (1) the noise of the lasers used and (2) the clock and data recovery circuit in the FPGA

As we add more users into the network, the noise floor begins to rise. This decreases the signal to noise ratio of the desired user and contributes to the higher bit error rate.

The high speed serial interface of the VirtexII-Pro is a hard IP placed on the FPGA programmable fabric and is therefore not itself programmable. It is designed to receive data with a high degree of transitions. Since we cannot guarantee that the received signal (aggregate of all transmitters) conforms to this specification, there are instances where errors are caused by failure of the clock and data recovery circuit of the high speed serial interface. We are currently working to bypass this circuit on the FPGA so that the incoming data can be clocked with an externally supplied clock

6. CONCLUSIONS

Our demonstration system proves the feasibility of uncoordinated

multiple access in optical networks. The building of the demonstrator combines design innovations in channel coding, FPGA logic implementation, and optical networks. The design of NL-TCM codes combined with IDMA provides a channel coding solution to optical uncoordinated multiple access. The Reed Solomon code ensures a network BER of less than 1e-9. To implement the algorithms on Virtex-II Pro FPGAs at optical rates, parallelization and pipelining techniques were extensively used. In the optical network, careful adjustment of the laser frequencies were necessary at the transmitter end. In the receiver, a high speed signal thresholding and retiming device was used to convert multilevel analog signals back to a digital signal. Though our goal was to demonstrate a 12 Gbps network with 6 transmitters, we were able to demonstrate a network of 7 Gbps which is an order of magnitude above that of current Ethernet technology.

ACKNOWLEDGEMENTS

This work was supported by the Space and Naval Warfare Systems Center – San Diego under contract No. N66001-02-1-8938.

REFERENCES

- [1] F.E. Ross and J.R. Hamstra, "Forging FDDI", IEEE Journal on Selected Areas in Communications, vol. 11, issue 2, Feb. 1993, pp. 181-190.
- [2] D. R. Boggs, J.C. Mogul, and C.A. Kent, "Measured Capacity of an Ethernet: Myths and Reality," ACM SIGCOMM '88 Symposium on Communications Architectures and Protocols, pp. 222-234.
- [3] L. Ping, K. Y. Wu, and L. Liu, "A Simple, Unified Approach to Nearly Optimal Multiuser Detection and Space-time Coding," ITW 2002, India, October 2002.
- [4] L. Ping, L. Liu, and W. K. Leung, "A Simple Approach to Near-Optimal Multiuser Detection: Interleaver-Division Multiple Access," IEEE Wireless Communications and Networking Conference, pp. 391-396, 2003.
- [5] G. Ungerboeck, "Channel Coding with Multilevel Phase Signals," IEEE Transactions on Information Theory, volume 28, pp. 52-67, 1982.
- [6] M. Griot, A. Vila Casado, W.Y. Weng, H. Chan, J. Basak, and R. Wesel, "Non-Linear Trellis Codes for Optical Applications," submitted IEEE International Conference on Communications.
- [7] Virtex-II Pro Platform FPGA Handbook (UG012) found at <http://www.xilinx.com/bvdocs/userguides/ug012.pdf>
- [8] D. Knuth, *The Art of Computer Programming Vol. 3 Sorting and Searching*, Addison-Wesley, Reading, MA, 1973, pp.229-232.
- [9] based on code found at <http://www.humanistic.org/~hendrik/>
- [10] T. Okoshi and K. Kikuchi, *Coherent Optical Fiber Communications*, Springer, 1988.
- [11] H. Al-Raweshidy and S. Komaki, *Radio Over Fiber Technologies for Mobile Communication Networks*, Artech House Publishers, 2002.
- [12] B. Saleh and M. Teich, *Fundamentals of Photonics*, John Wiley & Sons, 1991.