IEEE LTIMC 2004 – Lightwave Technologies in
Instrumentation & Measurement Conference
Palisades, New York, USA, 19 – 20 October 2004

Test setup for optical burst-mode receivers

Julien Faucher, Madeleine Mony, and David V. Plant

McGill University, Montreal, Quebec, H3A 2A7, Canada

***Abstract* — This paper proposes a test setup that can accurately measure the lock acquisition time of conventional and burst mode receivers subjected to phase, frequency, or amplitude steps. This work will be useful to network architects and circuit designers who wish to know with precision how many bits a receiver needs to acquire lock. The technique for measuring lock acquisition time, or preamble length, is based on an acceptable BER for the payload. Measurement results for the lock acquisition time of two commercially available OC-48 receiver chips will be presented.**

***Index Terms* — Burst switching, optical receivers, packet switching, phase detection, phase locked loops.**

I. Introduction

With the promise of delivering bandwidths in excess of 1 Gbps to the customer premises, optical multi-access networks such as passive optical networks (PONs) are one solution to the bandwidth hungry mix of voice, video, and data. An important component of such networks is a burst-mode receiver (BMR), which typically sits in the optical line terminal. BMRs are also an important component of optical burst- and packet-switched networks. The main goal of this paper is to present a test setup that can measure the lock acquisition time of BMRs subjected to phase, frequency, or amplitude steps.

The lock acquisition time of a BMR is an important performance characteristic because it dictates the number of preamble bits that should be reserved for clock and amplitude recovery. In order to increase the transmission efficiency, it is essential to minimize the preamble length. BMRs, unlike conventional receivers, have to handle packets of varying phase, amplitude, and possibly frequency. Hence, a BMR may have to realign the clock at the arrival of any new packet. This process must be quick because packets are typically short in multi-access networks. A conventional receiver requires a few thousand bits to acquire lock, whereas a BMR should require only a few bits. Throughout the paper, lock acquisition time will refer to the number of preamble bits that can guarantee a predetermined bit error rate (BER) on the payload. A precise measure of the lock acquisition time will be useful to network and protocol architects who wish to use the minimum number of bits possible in the preamble. Circuit designers wishing to measure the lock acquisition time of a clock and data recovery (CDR) circuit will also benefit from this research. An in-depth discussion of BMRs is beyond the scope of this text; the interested reader is referred to [1][3][7]-[10] for more information.

One technique for measuring lock acquisition time consists in monitoring the controlling voltage of the voltage-controlled-oscillator (VCO), which is the main block of a PLL-based CDR. Acquisition is considered complete when the envelope of the VCO control voltage stays within an acceptable percentage offset (usually taken as 2-5%) of the steady-state value. This technique has two drawbacks. First, it overestimates the lock acquisition time. It is not necessary for the clock to be perfectly aligned with the data before the payload becomes valid. Increasing the acceptable percentage offset will decrease the measured lock acquisition time, but the question then becomes "What percentage offset should be used to decide on the lock acquisition time?" The problem is that there is no direct relationship between the BER on the payload and the acceptable percentage offset. The second drawback of this technique is that the VCO control voltage is not always accessible on a package pin, especially on commercial CDR chips.

Given the limitations of using the VCO control signal to determine the lock acquisition time, an alternative technique is proposed in this paper. This alternative technique uses a decision rule based on an acceptable BER on the payload. Section II describes the implementation of a burst-mode test setup that uses this new technique, while Section III presents measurement results on two commercially available OC-48 CDR chips. Although of the conventional type, the receivers were nevertheless useful to debug the burst-mode test setup. Section IV is a concluding section.

0-7803-8722-8/04/$20.00 ©2004 IEEE

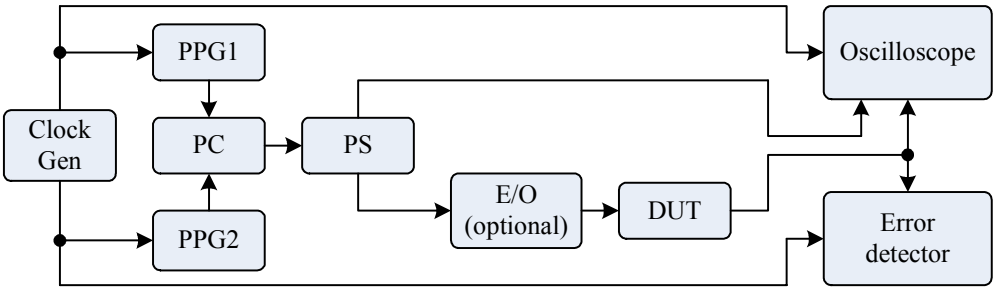

Fig. 1 Burst-mode test setup (the Matlab controller is not shown). PPG: pulse pattern generator, PC: power combiner, PS: power splitter.

II. Burst-Mode Test Setup Implementation

Conventional test equipment lacks some of the features that would facilitate the measurement of the lock acquisition time of CDR circuits. The test setup described herein (see Fig. 1) is a solution to this problem. It adds the following features to a bit error rate tester (BERT): 1) generation of two alternating packets with variable phase, frequency, and amplitude to emulate PON or bursty traffic, 2) measurement of BER vs. phase, frequency, or amplitude steps, and 3) measurement of the preamble length in order to guarantee a certain BER on the payload.

A. Packet generation

In order to test the lock acquisition time of a receiver against packets that vary in phase, frequency, and amplitude, one must first be able to generate such packets. This can be a challenging task in itself. Some pulse pattern generators (PPGs) offer the possibility of generating two alternating packets, but the phase, frequency, and amplitude of each packet cannot be set independently. A number of researchers have proposed solutions to overcome this problem [1]-[6]. Most solutions use a mix of programmable delay lines (PDLs), fiber delay lines (FDLs), fast optical switches, fast tunable transmitters, and variable optical attenuators (VOAs) to generate phase and amplitude steps. There is no mention of the possibility of generating frequency steps. A fast reconfigurable device is the workhorse of most solutions proposed thus far. Fast reconfigurable devices, such as fast optical switches and tunable transmitters, have a reconfiguration time that should be taken into account in the measurement of the lock acquisition time. In an attempt to isolate the problem of measuring lock acquisition time from the problem of measuring reconfiguration time, fast reconfigurable devices were avoided in the proposed packet generator. The number of uncertainties around the device under test (DUT) are therefore reduced.

As shown in Fig. 1, two PPGs and a power combiner are used for packet generation. No PDLs, FDLs, fast optical switches, or fast tunable transmitters were used in the setup. The phase, amplitude, and frequency of each packet can be set independently. The flexibility of this solution comes at the cost of synchronization complexity between the two PPGs. By sharing one common global clock, the two PPGs can achieve at best bit synchronization. They may therefore end up transmitting overlapping bits. Packet synchronization had to be implemented in order to emulate PON and OPS network traffic. In order to generate two non-overlapping packets, a Matlab GUI monitors and controls the two PPGs through GPIB. When one PPG is transmitting, the other is set to send 0's such that it appears to be silent – NRZ data is assumed. Assuming the two PPGs are generating data at the same frequency, the pattern loaded by the Matlab controller in each PPG has to be twice the length of each individual packet.

The performance specifications of the packet generator are determined by the two PPGs (an HP80000 and an Anritsu MP1763B). The maximum delay range between the two packets is 5 ns, while the delay resolution is 1 ps. The delay range sets a lower limit for the bit rate of the generated packets. For testing convenience, the delay range should be large enough to cover a full bit period. A 5 ns delay range covers the full bit period of a 200 Mbps signal. The delay resolution sets a higher limit on the bit rate of the generated packets. A 1 ps delay resolution is considered good because it reprensents 1 % of the bit period of a 10 Gbps signal. The packet generator can generate packets at up to 12.5 Gbps. In order to test the effect of amplitude steps on lock acquisition time, the amplitude of each packet can be set independently in the range from 0.2 V to 2 V.

B. How to measure the BER

The key performance measurement of communication links is the BER. Although BER measurements on

continuous data are common, such measurements on burst-mode data are more difficult. The most common way of testing a CDR is to use the recovered clock as the sampling clock for the error detector (ED). It will be argued next that a global or reference clock should be used instead.

Before the ED can start measuring a BER, it needs to synchronize its reference pattern to the incoming data pattern. The reference pattern can either be generated in hardware by the ED (for PRBS test patterns) or stored in internal memory (for user defined patterns). Synchronization times can range from μs to minutes and depend on the pattern length, the bit rate, and the synchronization method. Since the minimum synchronization time is in the μs timescale, it is important for the ED to not lose synchronization from packet to packet. This is especially true for short packets. SONET data, because it is continuous, will typically not cause synchronization losses. Synchronization is acquired when the link is first established and is never lost over the duration of the BER measurement. In contrast, synchronization may be lost with burst-mode data if the ED is clocked the usual way, i.e. with the clock recovered by the CDR. Two events can cause the ED to lose synchronization when operating with burst-mode data: 1) the CDR loses lock, and 2) the data signal was shifted with respect to the recovered clock signal, causing the incoming pattern to no longer correspond to the reference pattern programmed in the ED. Cases 1 and 2 may occur when the CDR is subjected to a frequency or amplitude step. Case 2 will also happen following a phase step. A phase step may cause the recovered clock to sample the incoming bits at the edge of the bit window. Under this condition, the CDR may output a bit either exactly when the ED expects it, or one bit ahead of time, or one bit too late. The situation that occurs will be random over the first few bits of the packet.

Due to the difficulty in maintaining synchronization, it was concluded that the test setup typically used for characterizing conventional receivers was not suitable for the characterization of BMRs. The main problem is that the signal under test, the recovered clock, is also used as a sampling clock. One solution is to use a reference clock as the sampling clock. If the sampling clock is derived from the same clock that is used to generate the data, then the ED should no longer lose synchronization because of a loss-of-lock or a bit shift.

Before BER measurements can be performed, the ED has to be loaded with a properly shifted version of one of the two packets. As specified earlier, the length of the reference pattern is twice the length of each individual packets if one common frequency is assumed. One of the two packets, packet 1, is used as a dummy packet whose purpose is to let the CDR reach steady-state before a phase, frequency, or amplitude step is applied. Any error that relates to packet 1 is ignored by masking the corresponding bit. BER measurements are therefore performed over packet 2 only. As mentioned earlier, a Matlab GUI applies circular shifts to each PPGs in order to make sure that packets 1 and 2 do not overlap. These circular shifts were used to derive the position of packet 2 in the reference pattern to be loaded in the ED. The Matlab GUI uploads the reference pattern to the ED and applies circular shifts.

The proposed methodology for making BER measurements on burst-mode data solves the loss-of-lock problem with the CDR. The ED is able to measure a BER even if the CDR loses lock temporarily after a frequency or amplitude step. Most of the errors occur while the CDR reacquires lock. The BER is therefore an indication of how quickly the CDR can lock and will help in the measurement of lock acquisition time (see Section C). The proposed methodology also solves the problem of the data signal being shifted with respect to the recovered clock. The ED is able to measure a BER even if bits appear one bit too early or one bit too late during the phase recovery process. After the phase has been recovered, bits will arrive at the ED exactly when they are expected no matter what happened during the phase recovery process. Measurements of the BER vs. phase step and amplitude step will be presented in Section III.

C. How to measure the preamble length

Measuring the number of preamble bits that can guarantee a predetermined BER on the payload was the ultimate goal of this research. Equipped with a tool to measure BER vs. phase, frequency, and amplitude steps, the preamble length can be measured using the block window feature of the ED – an Anritsu MP1764A. As argued in Section A, packet 1 can be though of as a dummy packet and its bits are therefore always masked by the ED. All measurements are performed on packet 2. The goal is to mask the minimum number of bits at the beginning of packet 2 to obtain a BER of 10^{-10} or better on the unmasked bits. The masked bits then correspond to the preamble length while the unmasked bits correspond to the payload. The block window patterns that are attempted should be chosen carefully in order to reduce testing time. It is important to note that the Anritsu MP1764A ED allows for the masking of bits on a 32-bit resolution.

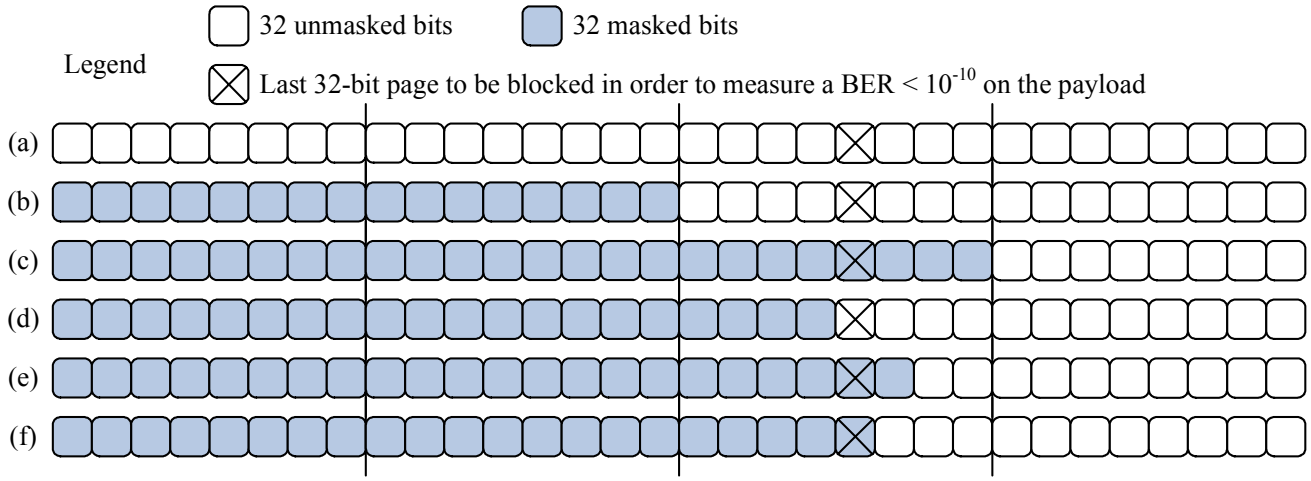

Fig. 2 Preamble length measurement using the non linear algorithm. The preamble length in this example is 21 × 32 = 672 bits.

The first implemented algorithm to find the preamble length was linear. Consider Fig. 2 as a case study. The packet length is 1024 bits. Only the packet of interest (packet 2) is shown. The gray and white squares represent 32 masked and unmasked bits, respectively. In this example, the 21st square, marked with an "X", represents the last 32-bit page that should be masked in order to guarantee a BER of 10^{-10} on the payload. In order to find this 32-bit page in an automated way, the linear algorithm considers 22 block window patterns (this includes the initial empty block window). Using the linear algorithm, 22 BER measurements are required to determine the preamble length for a given phase, frequency, or amplitude step. The complete linear algorithm took approximately two minutes to execute. While two minutes may not sound like much, it is important to keep in mind that the preamble length has to be measured many times for generating a plot of preamble length vs. phase, frequency, or amplitude step. Generating a plot of preamble length vs. phase step on a 10 ps resolution and a 1000 ps delay range (100 points) required about 3.3 hours. The linear algorithm is therefore inefficient and slow.

The second algorithm that was considered was non-linear. It required far less trials (6 instead of 22) and a preamble length could be found in about 20 to 30 seconds. Referring to Fig. 2, the search for preamble length stops when the incremental block window is 32 bits. More generally, the number of BER measurements required to determine the preamble length can be predicted using the following formula:

$$Nb.\,BER\ measurements = \log_2 N - 4 \qquad (1)$$

where N is the packet length. In the example of Fig. 2, N equals 1024 and the number of BER measurements equals 6. Generating a plot of preamble length vs. phase step on a 10 ps resolution and a 1000 ps delay range (100 points) required about 45 minutes. It should be noted that the linear algorithm will perform better in situations where the preamble length is short. This is because the number of BER measurements for the non linear algorithm is fixed at 6, while it varies for the linear algorithm.

III. Measurement Results

The burst-mode test setup described in the previous section was used to measure the BER at the output of two SONET receivers subjected to phase and amplitude steps. The number of preamble bits required to generate a 10^{-10} BER on the payload was also measured. Plots of BER vs. frequency step were not generated because both receivers support OC data rates only. The burst-mode test setup described in Section II would nevertheless be able to make such a measurement on continuous multirate CDRs.

A. Packet generation

The top waveform of Fig. 3 shows the eye diagram at the output of the packet generator. The two packets have equal frequency and amplitude, but a 160 ps phase difference. To test a receiver in continuous mode, the delay of one of the two PPGs can be adjusted such that the two eye diagrams appear perfectly superimposed. The bottom waveform of Fig. 3 shows the eye diagram at the output of the ADN2819 CDR from Analog Devices. Two eye diagrams corresponding to packets 1 and 2 can still be distinguished. What appears to be jitter really is due to the CDR that realigns the recovered clock with the data at the

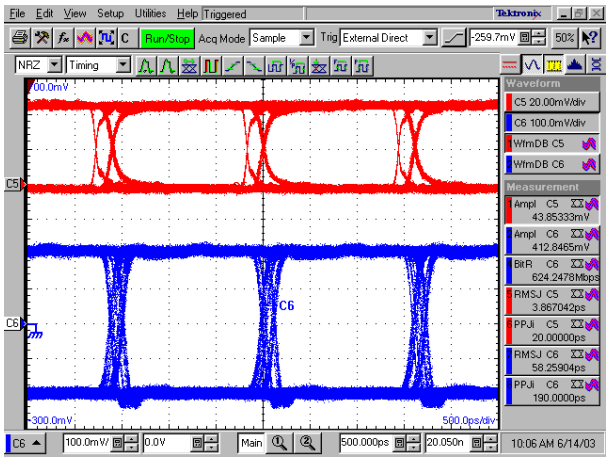

Fig. 3 Example of packet generation with a 160 ps phase step. A 2^{11} PRBS at 622.08 Mbps was used.

beginning of every packet. The horizontal opening of the eye diagram closes as the phase step approches π rad. For a π rad phase step, the first few bits of packet 2 are sampled right on the edge of the bit window. This worst case scenario generates many errors at the output of the CDR. In Section B, it will be shown that the measured BER is well above 10^{-10} for certain phase steps. Ideally, a BMR should produce a 10^{-10} or better BER for any phase step.

B. BER vs. phase step and lock acquisition time

Fig. 4(a) shows the measurements performed on a CDR chip from Analog Devices (part # ADN2819). The chip can be configured to operate at various data rates (OC-3, OC-12, OC-48, and GbE). Fig. 4(a) is a plot of the BER and preamble length vs. phase step. The packets are 2048 bits long (2^{11} PRBS), have a signaling rate of 622.08 Mbps, and 50 mV p-p of amplitude. Only packet 2 data is shown in the figure. For the worst case phase step between the packets, the CDR chip needs about 1400 bits to acquire phase lock. This is more than the number of bits in a 1 µs packet at OC-12 (~ 622 bits) and it shows that SONET receivers are not suitable for burst-mode applications.

Fig. 4(b) shows the measurements performed on a second CDR chip, this one from Maxim-IC (part # MAX3873). The chip can be configured to operate at 2.5 or 2.7 Gbps. Fig. 4(b) is a plot of the BER and preamble length vs. phase step. Both packets are 2048 bits long (2^{11} PRBS) and have an amplitude of 50 mV p-p. The signaling rates of packets 1 and 2 are 622.08 Mbps and 2.48832 Gbps, respectively. Only packet 2 data is shown in the figure. In the worst case, the CDR chip needs about 400 bits to acquire phase lock. This represents a significant portion (~ 16 %) of a 1 µs packet at OC-48. This shows once more that SONET receivers are not suitable for burst mode applications.

C. BER vs. amplitude step

Fig. 4(c) shows a plot of the BER vs. phase step for various packet amplitude steps. The ADN2819 device was used for these measurements. PPG1 generates a constant voltage amplitude of 2 V (50 mV at the CDR) while the amplitude generated by PPG2 is varied. The BER gets worse as the amplitude of packet 2 is decreased. At 0.77 V, the BER is higher than 10^{-10} for any phase step. This shows the importance of fast amplitude recovery. Fast clock recovery alone is not sufficient for burst-mode applications. There are three explanations for the poor performance of the receiver in Fig. 4(c). First, the amplitude of 0.77 V at the output of PPG2 corresponds to an amplitude of 19.25 mV at the input of the CDR. The worst-case differential input sensitivity of the CDR chip from Analog Devices is 10 mV p-p. This corresponds to a single-ended input sensitivity of 20 mV. A voltage amplitude close to the sensitivity of the receiver could therefore explain the poor performance observed. Second, the limiting amplifier in front of the CDR does not perform automatic threshold control (ATC). An ATC circuit with a short time constant is particularly important for BMRs, especially when a low amplitude packet follows a high amplitude packet. ATC is not as critical for SONET receivers because the amplitude of the incoming packets is relatively constant over time. Third, the DC rejection provided by AC coupling can cause baseline wandering. This problem is avoided in SONET networks by using data scrambling. Data scrambling limits the number of consecutive identical digits (CIDs), which are one source of baseline wandering. DC coupling is preferred over AC coupling for BMRs in order to speed up ATC and to avoid baseline wandering.

VII. Conclusion

This paper has presented a test setup that can accurately measure the lock acquisition time of conventional and burst-mode receivers subjected to phase, frequency, or amplitude steps. The burst-mode test setup was used to make measurements on two commercially available CDR chips. The setup offers numerous features that are useful for testing BMRs: 1) Generation of two alternating packets; the phase, frequency and amplitude can be set independently for each packet. 2) Good delay range [5 ns] and delay resolution [1 ps] between packets. 3) Does not use fast optical switches and tunable transmitters because the reconfiguration time of these devices introduces

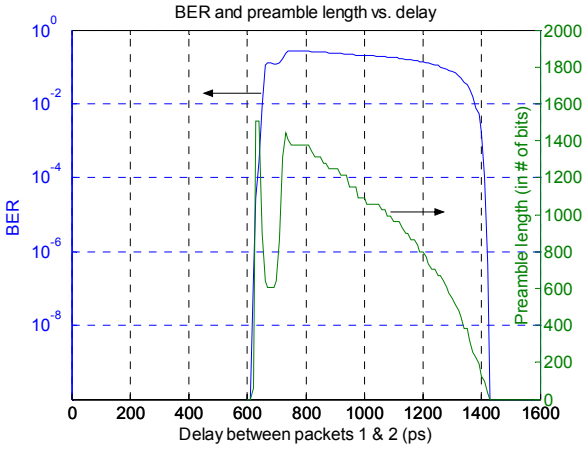

(a) AND2819, 2^{11} PRBS at 622.08 Mbps.

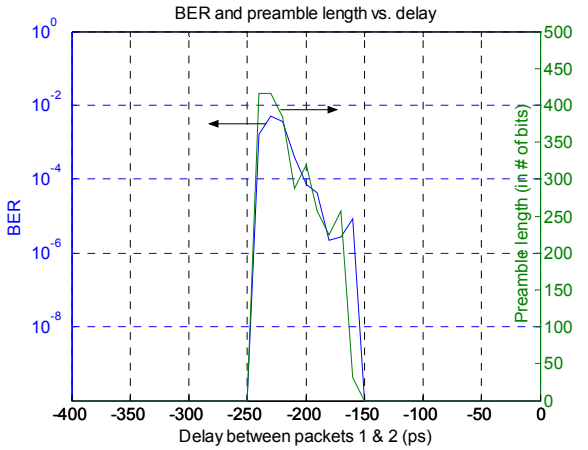

(b) MAX3873, 2^{11} PRBS at 2.48832 Gbps.

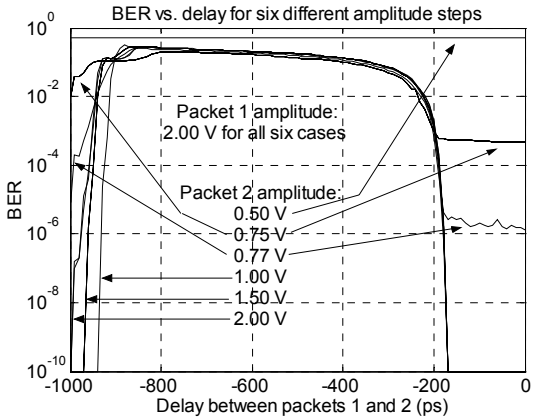

(c) AND2819, 2^{11} PRBS at 622.08 Mbps.

Fig. 4 (a)-(b) BER vs. phase step and preamble length vs. phase step. (c) BER vs. amplitude step.

uncertainties in the measurement of the lock acquisition time. 4) Measurement of the BER vs. phase, frequency, and amplitude steps. 5) Measurement of the preamble length (lock acquisition time) of a receiver with a guaranteed predetermined BER on the payload. 6) Automated measurements can be performed using a Matlab GUI to control the burst-mode test setup.

ACKNOWLEDGEMENT

This work was supported by the Natural Sciences and Engineering Research Council of Canada (NSERC), the Fonds québécois de la recherche sur la nature et les technologies (formerly FCAR), McGill University (Tomlinson Fellowship), and industrial and government partners, through the Agile All-Photonic Networks (AAPN) Research Network.

REFERENCES

[1] S. Rubin, E. Buimovich, G. Ingber, and D. Sadot, "Implementation of ultrafast widely-tunable burst-mode 10 Gbit/s transceiver," Electron. Lett. , vol. 38, pp. 1462 -1463, Nov. 2002.

[2] Y. Yamada, M. Togashi, K. Habara, and T. Matsunaga, "Burst-mode differential receiver for optical packet communication," Electron. Lett., vol. 32, pp. 1500-1501, Aug. 1996.

[3] H. Nishizawa, Y. Yamada, K. Habara, and T. Ohyama, "Design of a 10-Gb/s burst-mode optical packet receiver module and its demonstration in a WDM optical switching network," J. Lightwave Technol., vol. 20, Jul. 2002.

[4] C. A. Eldering, F. Herrerias, R. Martin, P. J. Garcia, M. Martin, and J. C. Coedo, "Amplitude and phase recovery systems for burst-mode fiber optic receivers," Proceedings of SPIE, vol. 1975, Bellingham, WA, 1993, pp. 292-298.

[5] N. Pleros, K. Vyrsokinos, C. Bintjas, K. Yiannopoulos, K. Vlachos, H. Avramopoulos, and G. Guekos, "All-optical clock recovery from short asynchronous data packets at 10 Gb/s," IEEE Photon. Technol. Lett., vol. 15, no. 9, pp. 1291-1293, Sept. 2003.

[6] Y. Yamada, S. Mino, and K. Habara, "Ultrafast clock recovery for burst-mode optical packet communication," in OFC Tech. Dig., San Diego, CA, Feb. 1999, paper TuI5, pp. 111–113.

[7] Y. Ota, and R. G. Swartz, "Burst-mode compatible optical receiver with a large dynamic range," J. Lightwave Technol., vol. 8, pp. 1897 -1903, Dec. 1990.

[8] H. Tagami, S. Sato, M.Nogami, Y. Matsumoto, K. Motoshima, and T. Kitayama, "156 Mbit/s burst-mode transceiver with a new bit synchronization technique for PON application," 22nd European Conference on Optical Communication (ECOC '96), vol. 4, pp. 97-100, Sept. 1996.

[9] L. Lunardi, S. Chandrasekhar, R.G. Swartz, R.A. Hamm, and G.J. Qua, "A high speed burst mode optoelectronic integrated circuit photoreceiver using InP/InGaAs HBT's," IEEE Photon. Tech. Lett., vol. 6, pp. 817-818, Jul. 1994.

[10] Y. Yamada, S. Mino, and K. Habara, "Ultra-fast clock recovery for burst-mode optical packet communication," Technical Digest of the Optical Fiber Communication Conference, 1999, and the International Conference on Integrated Optics and Optical Fiber Communication (OFC/IOOC '99), vol. 1, pp. 114-116, 1999.