RC23067 (W0401-087) January 19, 2004
Electrical Engineering

IBM Research Report

120 Gb/s VCSEL-Based Parallel Optical Interconnect and Custom 120 Gb/s Testing Station

Daniel Kuchta, Young Kwark, Christian Schuster, Christian Baks, Chuck Haymes, Jeremy Schaub, Petar Pepeljugoski, Lei Shan, Richard John, Daniel Kucharski, Dennis Rogers, Mark Ritter
IBM Research Division
Thomas J. Watson Research Center
P.O. Box 218
Yorktown Heights, NY 10598

Jack Jewell, Luke Graham
Picolight Inc.
Boulder, CO 80501

Karl Schrödinger
Infineon Technologies AG
Fiber Optics
D-13623 Berlin
Germany

Alexander Schild, Hans-Martin Rein
Ruhr-University Bochum
D-44780 Bochum
Germany

Research Division
Almaden - Austin - Beijing - Haifa - India - T. J. Watson - Tokyo - Zurich

120 Gb/s VCSEL-Based Parallel Optical Interconnect and Custom 120 Gb/s Testing Station

Daniel Kuchta, *Senior Member IEEE,* Young Kwark, *Member IEEE*, Christian Schuster, *Member IEEE,* Christian Baks, Chuck Haymes, Jeremy Schaub, *Member IEEE*, Petar Pepeljugoski, *Senior Member IEEE,* Lei Shan, Richard John, Daniel Kucharski, *Student Member IEEE,* Dennis Rogers, *Member IEEE*, Mark Ritter, Jack Jewell, *Member IEEE,* Luke Graham, Karl Schrödinger, Alexander Schild, *Student Member IEEE,* and Hans-Martin Rein, *Fellow IEEE*

***Abstract*—A 120Gb/s optical link (12 channels at 10 Gb/s/ch, both a transmitter and an receiver) has been demonstrated. The link operated at a BER of less than 10^{-12} with all channels operating and with a total fiber length of 316 m which comprises 300m of next generation (OM-3) multimode fiber (MMF) plus 16m of standard grade MMF. This is the first time that a parallel link with this bandwidth at this per channel rate has ever been demonstrated. For the transmitter, a SiGe laser driver was combined with a GaAs VCSEL array. For the receiver, the signal from a GaAs photodiode array was amplified by a 12 channel SiGe receiver IC. Key to the demonstration were several custom testing tools, most notably a 12-channel pattern generator. The package is very similar to the commercial parallel modules that are available today, but the per-channel bit rate is three times higher than for the commercial modules. The new modules demonstrate the possibility of extending the parallel optical module technology that is available today into a distance*bandwidth regime that is unattainable for copper cables.**

***Index Terms*— Crosstalk, Driver circuits, Optical interconnections, Optoelectronic packaging, Parallel optical link, Semiconductor lasers, Receivers, Testing.**

I. Introduction

EXISTING parallel connections between computing nodes in servers have already reached an aggregate transfer rate in the 30Gbaud range [1]. Future generations of servers are likely to have point to point connection bandwidths above 100Gbaud [2]. At these connection speeds, the choice between parallel copper or parallel fiber is not so obvious, with one important issue being the packaging and testing of the components. This paper describes the development of packaging and testing tools for a high speed parallel fiber link intended for short distance communication between computing nodes. The design and testing of this link have been performed in a cost conscientious manner so as to create a technology which can compete with copper economically, as optics already outperforms copper in speed (at intercard distances and greater) and physical size.

There are numerous papers on 10 Gb/s serial VCSELs [3], [4] and since the ratification of the 10GbE standard [5], many of these are now in products [6]. There are also reports on parallel links [7], and 3 Gb/s/ch parallel links are at the product level [8]. At the 10 Gb/s/ch level there has been a twelve channel link proposal [9] and the demonstration of a 12 channel receiver IC [10],[11]. There is prior work on four channel links at 10 Gb/s/ch using edge emitting lasers [12],[13], VCSELs [14],[15], and coarse wavelength division multiplexed VCSELs on a single fiber [16]. Kurata's work [14] is notable as it is a full duplex link at 10 Gb/s/ch whereas others, including this paper, report only simplex link operation. There are also products which have achieved an aggregate transfer rate above 100 Gb/s [17] but at the expense of large fiber count. Recently [18] the technical details of these high fiber count modules have appeared in the literature. High fiber count solutions are not expected to be cost competitive with copper until the price of high fiber count connectors and jumper cables are significantly reduced.

The paper is organized as follows. Section II covers the SiGe driver, receiver IC and the VCSEL and PIN arrays. In Section III the details of the flex based package and ball grid array (BGA) connector are discussed. Section IV introduces the low cost custom tester used to generate twelve channels of data for simultaneous testing. The last section, Section V, summarizes the testing results obtained.

Manuscript received January 15, 2004. This work was partially supported by NIST through the Photonics CAD ATP and by the German Ministry of Research (BMFT).

D. Kuchta, Y, Kwark, C. Schuster, C. Baks, C. Haymes, J. Schaub, D. Rogers, M. Ritter, L. Shan, R. John, and P. Pepeljugoski are with the IBM T. J. Watson Research Center, Yorktown Heights NY 10598 USA (phone: 914-945-1531; fax: 914-945-4134; e-mail: kuchta@us.ibm.com).

D. Kucharski was with the IBM T. J. Watson Research Center, Yorktown Heights NY 10598 He is presently with Cornell University, School of ECE, Ithaca NY 14853 USA

L. Graham and J. Jewell are with Picolight Inc., Boulder CO 80501 USA

K. Schrödinger is with Infineon Technologies AG, Fiber Optics, D-13623 Berlin, Germany

A. Schild and H. -M. Rein are with Ruhr-University Bochum, D-44780 Bochum, Germany

II. Driver and Receiver Circuits

A. SiGe Driver Array

The driver IC is a monolithic 12-channel laser diode driver (LDD) array fabricated in IBM's SiGe BiCMOS technology with f_T = 120 GHz [19]. A simplified schematic diagram of a single LDD channel is shown in Fig. 1. It consists of a limiting preamplifier with a 100 Ω floating input termination followed by a transconductance amplifier (TCA) output stage. The preamplifier uses a Cherry-Hooper gain stage [20]. It improves transition times and assures a constant signal amplitude at the input to the TCA. The minimum input voltage is 150 mV_{p-p} with a common mode voltage of 1.25 V, which is compatible with the LVDS standard [21]. The inputs to each of the channels can be either DC or AC coupled. The signal path is fully differential to reduce crosstalk except for the output stage, which has a single-ended current output to accommodate a common-cathode laser array. However, the output of each channel has an adjacent ground pad that provides a low-inductance return path for the modulation current.

A high output impedance TCA is well suited for current modulation, since it can tolerate large variations in laser's series resistance, up to 140 Ω in this case. In addition, the output of the TCA is DC coupled, which is beneficial for array applications, as it eliminates the need for large coupling capacitors and separate laser bias circuitry. The TCA consists of a current switching differential pair Q1–Q2 and a PFET current source M2 as shown in Fig. 1. Modulation is accomplished by diverting a part of M2 drain current through Q2, while the remaining current flows through the VCSEL. Emitter degeneration resistors R_E improve linearity and are bypassed with capacitor CE to speed up output transitions. Bias currents I_H and I_L can be adjusted off-chip to set the desired modulation levels. All 12 channels share the same modulation and bias controls. The output modulation current is adjustable up to 10 mA_{p-p} and the DC bias is adjustable up to 7 mA.

The LDD consumes 100 mW per channel (including the VCSEL) from a single 3.3 V power supply. At the die level, the 20–80% output transition times are less than 20 ps measured with a 50 Ω load. The driver operates in an open-loop configuration with no optical feedback. However, it has an integrated temperature reference circuit, which compensates for temperature variations in the LDD. The driver IC die size is 1.25 mm × 3.50 mm and is pad limited; however the extra chip real estate is useful for incorporating on-chip decoupling capacitance. The channels are on a 250 μm pitch to facilitate direct and uniform wire bonding to the VCSEL array.

B. SiGe Receiver Array

The 12-channel receiver array has been extensively described in [10] and shall, therefore, be recapitulated only briefly. It was fabricated in Infineon's SiGe bipolar technology B7HF, characterized by 0.35 μm lithography and about 72 GHz f_T [22]. The chip size is 2.8 mm × 4.6 mm, and the pitch of the channels is 250 μm.

In each channel, a high-gain transimpedance amplifier with limiting output stage transforms the input current, generated by the photodiode, to a differential output voltage v_Q with constant swing (2 × 250 mV_{p-p} = 500 mV_{p-p}). The specified input current swing ΔI_I ranges from 20 μA_{p-p} to 240 μA_{p-p}, resulting in a maximum (non-linear) transimpedance as high as 25 kΩ. A simplified block diagram of the transimpedance amplifier is shown in Fig. 2. The circuit is driven by the photodiode, which is directly bonded to the amplifier's differential input, and loaded by 50 Ω (e.g. terminated transmission lines). The required bias voltage for the photodiode of 1.24 V, which is sufficiently high for the initially specified pin photodiodes, is generated on-chip. Note, that here the circuit version for positive supply voltage (+5 V) is used (cf. Fig. 5 in [10]). Moreover, it is assumed that the succeeding clock and data recovery circuit has a separate supply voltage V_{cco}, which is usually lower than that of the amplifier array to reduce power consumption. In this case, DC coupling is simply obtained by connecting the on-chip output resistors (100 Ω) as well as the 50 Ω loading resistors to V_{cco}. The required decoupling capacitors are implemented on the chip. For the measurements presented below, AC. coupling is used at the output, since the 50 Ω input impedance of the oscilloscope is connected to ground.

The amplifier is composed of three similar amplifier cells, each consisting of cascaded differential transimpedance, emitter follower, and transadmittance stages. The output transadmittance stage (3rd cell) operates in the limiting mode even at minimum input current swing, thus working as a current switch. An automatic offset control (AOC) is required for each channel to compensate for the average value of the input signal current (see Fig. 2). The improvements by use of the input dummy capacitor C_{dum}, which roughly equals the photodiode capacitance C_{PD}, have been discussed in [10].

At the nominal supply voltage of +5 V, the power consumption is 92 mW per channel. It should be mentioned, that the main challenge of the design was to reduce crosstalk between the channels to a sufficiently low level. This aim has been achieved by applying adequate measures discussed in [10], as confirmed by detailed measurements.

C. VCSEL and PIN Photodiode Arrays

The VCSELs are 850 nm oxide confined devices with sub-milliamp threshold fabricated in a monolithic 12ch array on a 250 μm pitch. The array has a common cathode contact which is the substrate. The array is similar to ones used in 3 Gb/s production modules , but have been optimized for reliable 10 Gb/s performance over temperature. The series resistance of the VCSEL is 70 Ω. The VCSEL is multi-transverse mode with an RMS spectral width, 420 pm, within the specifications for 10-Gigabit Ethernet. With modulation current from 2.5 mA to 7.5 mA, the VCSELs are able to

maintain an open eye at 10 Gb/s over temperature.

The photodiodes are GaAs PINs fabricated in a monolithic 12 channel array on a semi-insulating substrate.. The photodiode capacitance is 238 fF at a reverse bias of 1.2 V and has a bandwidth of 7.0 GHz measured into a 50 Ω load. The dark current is 1.3 nA at -1.2 V.

III. PACKAGING

A. Optical Module

One of the constraints imposed on this project was to preserve as much of the existing (3 Gb/s) package and assembly infrastructure as possible. This required that the overall mechanical housing and the optical coupler and the location of the VCSEL within the housing remain unchanged to preserve the assembly and alignment procedure. The optical transmitter and receiver were packaged separately in a module that conforms to the SNAP12 multi-source agreement [23]. The outer dimensions of the package are 13 mm × 18 mm × 41 mm. The package accepts an MTP® connector at the optical interface and has a pluggable 100 pin ball grid array connector, the MEG-Array® [24], at the electrical interface. The module is pictured in Fig. 3. Within the module is a flexible printed circuit board (flex). The primary function of the flex is to perform a spatial transformation of the electrical signals from the nearly 1 mm pitch at the BGA connector to the 125 μm pitch at the input of the driver chip and to bend these signals 90 degrees out the plane of the system circuit board. The outer dimensions of the flexible printed circuit board within the housing remained constant but within those dimensions the wiring could be changed. The electrical connector was rotated by 90 degrees to reduce wiring skew in the differential channels; this had the salutary effect of preventing inadvertent mating with 3Gb/s testing fixtures since pin compatibility was not maintained..

B. Transmission Lines on Flex

The flex circuit is comprised of two dielectrics, a thin flexible layer and a thicker supporting layer. The thick supporting layer is 10 mils of FR4. The flexible layer is 2 mils of a polyimide material. The polyimide material has a loss tangent of less than 0.005 up to 20 GHz [25]. The thin flexible layer also offers the advantage of accommodating small line widths and spacings. The 100 Ω differential transmission lines which connected the MEG-Array® to the driver chip were realized with dimensions of 2 mil width and 2 mil space. With these small dimensions the transmission lines were able to be routed into the MEG-Array® pin field without any change to their dimension thus preserving their impedance throughout their entire length. The length of transmission line varied from channel to channel from 13.9 mm to 30.0 mm with an average length of 20 mm. The transmission lines were separated by 0.33 mm throughout most of their length to minimize crosstalk. Using a 4-port vector network analyzer, differential measurements made on straight sections of transmission lines showed a loss of 0.09 dB/mm at 15 GHz. This measured value is almost double the simulated value based on the conductor and dielectric losses. Part of the discrepancy is due to a slight overetch of the lines and the roughness of the copper-polyimide interface which was not taken into account. Additionally, the loss tangent in the GHz regime of the thin (0.5 mil) soldermask material may be higher than the quoted value of 0.04 obtained at 1 MHz. The -3 dB bandwidth of a straight 29 mm line is 18.8 GHz. The bandwidth of a 20 mm line with six 45 degree bends is 28 GHz It should be noted that these -3 dB frequencies do not represent corner frequencies of a single pole system since S21 is dominated by the conductor losses and falls off as the square root of frequency .

Each differential pair was separated from its nearest neighbor in the BGA connector by at least one row and column of pins. To the extent possible, the separation row and column were connected to ground to form a ground cage around the signals to provide some electromagnetic shielding. The transmission line lengths within a differential pair were matched to better than 100 μm, which corresponds to an intrachannel skew of less than 0.5 ps.

C. Driver and Receiver Mounting and Decoupling

Both the driver and receiver ICs are intended to be mounted within a cutout in the flex circuit to minimize the thermal impedance to the heatsink. The lowest cost method of producing the cutout in the flex is performed with a mechanical router. However, the tolerances on the cutout from mechanical routing are fairly large and this slack gets taken up by the wirebonds to the chip for both the signal and power supply connections. At 3 Gb/s the wirebond lengths are not a significant concern but at 10 Gb/s is it prudent to minimize the wirebond lengths. Several changes were made to the chip mounting to improve the signal integrity and power supply decoupling for 10 Gb/s operation. The first change was to switch from mechanical to laser routing of the cutout. This change reduced the tolerance on the cutout by a factor of 3 and reduced the overall size of the cutout since the laser based cutting can produce square corners. The second change was to place a ceramic “shim” beneath the chips. The ceramic shim has a metallic surface which serves as the ground plane for the chips. The shim elevates the height of the chips with respect to the flex surface. This has the effect of lengthening the wirebonds on the electrical side but shortens them on the optical (VCSEL or PIN) side. The electrical side is fully differential so an increase in wirebond length has a minor impact on the performance. The shim is also sized larger than the chips on the short edge to accommodate three 200 pF capacitors for power supply decoupling on each side. This arrangement allowed the decoupling capacitors to be placed as close as 70 μm from the chip edge and to be connected to the chip with a wirebond as short as 400 μm. The differential inputs to the driver chip each have a pair of wirebonds

averaging 660 μm in length.

The VCSELs and photodiodes are mounted on a Au plated copper submount which is attached to the flex by conductive epoxy. For the VCSELs this submount also serves as the common cathode contact. The edge of the VCSEL die is 750 μm from the edge of driver chip. Due to a height difference, the wirebond length from driver to VCSEL is 990 μm. The ground return for the output stage of the driver chip is attached to the flex ground through thirteen 730 μm long wirebonds. Fig. 4 shows the component placement and wirebonds within a completed transmitter assembly. Fig. 5 Photo of assembled receiver IC and photodiode array on flex circuit showing wirebonds and decoupling capacitors.

D. Meg-Array® Connector

The MEG-Array® connector is a mezzanine connector system with ball grid array terminations rated for 10 Gbps differential signaling [24]. Initial time domain reflectometry (TDR) of the MEG-Array® attached to the flex revealed a large capacitive discontinuity rather than the expected inductive discontinuity. This was analyzed in detail and presented in [26]. To summarize, the large size of the solder balls on this connector and the thin dielectric layer on the flex form a parasitic capacitance which is created between the solder pads and the ground plane of the flex. To reduce this parasitic, the ground plane was removed beneath the solder pads. In a second refinement, the solder pads and solder mask opening were reduced in size from the recommended 25 mil diameter down to 18.5 mil to further minimize the parasitic capacitance while remaining within the manufacturing assembly capabilities. Fig. 6 shows an eye diagram measured through the MEG-Array® at 10.3 Gb/s. In this measurement, the MEG-Array® is sandwiched between two test cards with approximately 1.5 inches of 100 Ω differential transmission line on either side. The test cards have the ground plane removed under the solder pads in a manner similar to that implemented on the flex. In this test, the input rise and fall times were 20 ps (20-80%). The output rise and fall times are 30 ps and the resulting eye is quite good for 10 Gb/s. It should be noted that there are no vias in the entire signal path.

E. Thermal Impedance

The thermal impedance of the package was a concern because of the extra layers added to the existing package by the ceramic shim and the conductive epoxies used to attach the IC's and the shim to the package. The close proximity of the driver and VCSEL would mean that the VCSEL's junction temperature would be elevated by the driver's power dissipation. To address this concern, detailed 3D thermal modeling of this package configuration was performed and compared with measurement [27]. Fig. 7 shows a comparison of the modeling and measurement results at two different ambient temperatures of 20° C and 30° C. The measured junction temperature of the VCSEL array has a slight parabolic shape and is actually hotter at the ends than in the center by a few degrees. A chip thermistor was added as one of the components on the flex and placed as close to the VCSEL as possible. Using this thermistor and the VCSELs wavelength shift with temperature, we were able to accurately determine that the driver chip rose 14° C above ambient while the VCSEL's junction rose by 28° C. When operated at room temperature (23° C) with an air flow of 8 ft/s, the interior of the package is at 45° C.

IV. Low Cost Testing

A. Possible Concepts

One of the most expensive aspects of developing parallel optical interconnects is the multi-channel test equipment needed to characterize them. At the time that the development of the technology presented here occurred, multi-channel systems capable of testing 12 channels at 10 Gb/s/ch were estimated to cost over 1 million US$. These commercially available solutions were also undesirable because they were not frequency agile and limited to testing in a narrow frequency range around 10 Gb/s. Investment in such a system for a fixed data rate is more suited for a manufacturing environment than for research and development. A cheaper alternative is to use a 12 channel redistribution box. A redistribution box accepts one pattern generator signal, amplifies it, passively splits it into 12 identical channels and amplifies it a final time. Several vendors offer this type of solution. An even cheaper solution, the one that was ultimately employed, is to build a tester from pseudo random bit sequence (PRBS) generator chips. A single chip solution for 12 channel testing has been alluded to in [28] but such chips dissipate a lot of power and can be difficult to package and use due to the cooling requirements. The 10 Gigabit Ethernet Standard [8] requires that all physical layer (PHY) IC's contain a PRBS generator and error detector. These chips, which are not frequency agile, are readily available at very low cost since they are produced in volume.

B. Multi-channel tester using PRBS generator ICs

A custom low cost multi-channel tester was thus realized by arranging 12 separate PHY chips into a tower which also holds the power supplies and cooling fans, and provides a platform for the device under test. This tester is shown in Fig. 8. It is a semi-circular tower with the optical module placed on an elevated test card in the center. Twelve test boards are arranged radially in the tower, each holding a 10 Gb/s Ethernet PHY chip. Fig. 9 shows a picture of the test card with a transmitter mounted in the center. This geometry of the tester allows equal lengths between the module and each of the twelve test boards. It also allows the cable length to be very short (4 inches), to minimize signal degradation. The short cables also result in small, < 6 ps, interchannel skew. The 10GbE PHY chips are used as pattern generators for transmitter testing and as error detectors for receiver testing. They can generate a number of test patterns including the standard 2^{31}-1 PRBS (PRBS31) and a clock÷22 square wave. The 10 Gb/s outputs are synthesized from 156 MHz reference

clock oscillators which eliminate the problem of distributing a 10 GHz clock to 12 channels. However, the jitter at the output is higher than what would be attained by using a PRBS generator chip driven by a high speed clock. In this particular implementation, each test channel has its own reference oscillator. Due to the slight frequency difference of 50 ppm between oscillators, each channel is operating asynchronously with respect to the others. Synchronous operation is also possible using a single reference oscillator distributed to each of the channels. The particular reference oscillator used has two outputs, the second of which is taken to a connector panel on the rear of the tower and is used to provide a trigger signal for making eye diagram measurements on an oscilloscope. Rise and fall times for each channel are < 30 ps (20-80%). The voltage output amplitude is adjustable from 500 mV to 2.0 V, and signaling is 100 Ω differential. Fig. 10 shows a differential eye measured at the output of one channel of the tester. Individual enabling and disabling of channels allows for crosstalk and single channel measurements. All test boards have a PC interface for remote control and data collection. The testing tower consumes only 140 W with all 12 boards enabled. The receive side of the testing tower has built in clock recovery and an error detector for the PRBS31 pattern. The error detector can count up to 88 errors before overflow which makes it suitable for link testing when BER < 8×10^{-9}. The tower can also be configured to operate at a data rate of 3 Gb/s, in addition to 10.3 Gb/s

C. *Test Card*

The test card, shown in Fig. 9, functions as an interposer card between the tester and the particular electrical connector for the parallel module under test. In this manner, different test cards can be designed to interface different parallel modules with the same tester. Since the test card is in the critical path it is important that it has the minimum impact possible. The test card was kept as small as possible to reduce the attenuation of the transmission lines. The small size drives the connector choice which in this case is the edge mounted GPPO™ connector. The design and measurement of the launch structure for the GPPO™ connector has been described in detail in [26]. To additionally minimize the total discontinuity from the test card, a flex printed circuit board technology is employed. It consists of a 3 mil polyimide layer on top of a 57 mil thick FR4 core. The thin polyimide layer is used to support 100 Ω differential transmission lines of 3 mil width and space. Similar to the optical package, this permits the transmission lines to enter the MEG-Array® pin field without changing their dimensions. The test card is less constrained than the flex in the optical package in that the transmission lines are free to enter the MEG-Array® pin field from all four sides and thus there is less potential for crosstalk. 100 nF capacitors are placed in the signal path near the GPPO™ connector to act as DC blocks.

To reduce attenuation, the transmission lines on the test card were not covered with solder mask. In doing this, the copper traces must be protected from corrosion. Au plating over Ni is often used for this purpose but Ni is ferromagnetic which will increase the line attenuation unless a very thick layer of Au is used. Ni plating also results in a fairly rough surface which increases the attenuation. In our implementation, Sn plating is used over the copper to protect the lines. Although Sn has lower conductivity than Au and should have higher skin loss, measurement on 1.4 inch test lines revealed that the attenuation of Sn plated lines was actually lower than Au/Ni plated lines by as much as 0.3 dB up to 17 GHz. Beyond 17 GHz, the two line variants were nearly identical. The bandwidth of a 1.4 inch test line (without connectors) with Sn plating was > 30 GHz. The lines on the test card vary from 1.29 inch to 1.69 inch with an average length of 1.46 inch. The intrachannel skew was controlled to less than 1.6 ps. The actual bandwidth on the test card is dominated by GPPO™ connector. A differential measurement of the test card starting at the GPPO™ connector going through the MEG-Array® and ending on the package flex with a differential probe at the driver chip site had a -3 dB bandwidth of 7.5 GHz**.** Again, this -3 dB frequency is not the corner frequency of a single pole system as the transmission only decreases with the square root of frequency.

D. *Additional Custom Equipment*

While the testing tower is useful as a 12 channel pattern generator and error detector, several other pieces of custom test equipment are needed to fully characterize a transmitter and receiver separately and also the whole link. One such piece is a single pole twelve throw (SP12T) microwave switch for switching the 12 trigger signals to an oscilloscope. A second item is a dual SP12T microwave switch with 50 Ohm terminations on the unused inputs. This is used for receiver characterization where one side of one channel's differential outputs is directed to an error detector, the other side is directed to an oscilloscope while the other 11 channels are differentially terminated. The measured -3 dB bandwidth of this dual switch was over 40 GHz.

A third item is a custom multimode fiber optic switch that is inserted in the link path. A schematic diagram of this switch is shown in Fig. 11. This switch permits any one of the 12 channels to be brought out (tapped) to a pair of connectors where additional optic test equipment can be attached. The other 11 channels pass through the switch. The switch has more than 12 positions so that the breakout channel can be moved to unused position allowing all 12 channels to pass through. In our particular implementation, the breakout channel is sent to a pair of 1×2 fiber optic switches. On the transmitter side, the breakout channel can be routed to a golden receiver (linear, low noise, high bandwidth receiver) for separate transmitter eye characterization and to the input of an optical attenuator for BER testing. On the receiver side the breakout channel accepts input from a golden transmitter (commercially available 12.5Gb/s 850nm VCSEL) for receive

channel characterization or from the output of the optical attenuator for BER characterization. Using the testing tower, the two microwave switches, and the optical switch, the transmitter, receiver and link can be completely characterized under computer control without changing the setup.

V. Results

A. Transmitter

The optical transmitter is first characterized separately for optical output properties and crosstalk. For all of these measurements, a short 2 m breakout cable is connected from the transmitter to the optical switch. A golden receiver is used for the AC measurements.

Fig. 12 shows 12 eye diagrams measured on a transmitter before the optical coupler was attached. All 12 eyes are wide open and quite uniform. The optical rise and fall times are 38 ps (20%-80%) measured with a 12 GHz receiver. With the optical coupler attached and the module fully assembled the average optical power is 1.1 mW with a +/-15% variation as shown in Fig. 13. The extinction ratio can be varied from 1.3 to > 4 through the use of one of the control voltages on the driver chip. At an average extinction ratio setting of 2.8 the variation across the array is +/-10% as shown in Fig. 14. While the coupled power measurement is affected by variations in the coupler and alignment tolerances, the extinction ratio is not. The uniformity of the extinction ratio is indicative of the uniformity of the VCSEL L-I curves and the driver's signal levels.

One of the concerns for the use of multimode fiber is the ability of the coupled VCSEL to conform to the encircled flux measurement [29]. The transmitter encircled flux for all 12 channels was measured and found to be compliant for use with OM-3 multimode fiber [30]. Fig. 15 shows the encircled flux data for all 12 channels. Open eyes were obtained at 300m fiber length using a 12 fiber ribbon of 50/125 OM-3 MMF fiber. At the end of 300 m, the rise and fall times were 44 ps (20%-80%). The eyes were measured at 500 m using a single strand of OM-3 MMF. Again, all 12 eyes were open with rise and fall times of 51 ps. At the end of a single 750 m strand of OM-3 MMF only 6 of the 12 eyes were open and the rise and fall times were 60 ps. Even though one of the 12 VCSEL's encircled flux was at the edge of the specification, the 300 m distance was easily met and exceeded.

Crosstalk within the transmitter is measured as follows. For amplitude crosstalk, a BER curve is measured on the channel under test using a golden receiver with all other 11 channels off. This becomes the reference curve for that channel. Next, each of the 11 aggressor channels is turned on one at a time and a BER curve of the channel under test plus one aggressor is measured. Finally all 12 channels are turned on simultaneously and the BER curve of the channel under test is measured. The 12 crosstalk BER curves are compared to the reference curve and the difference is the crosstalk penalty in dB of optical power. In this manner, 169 BER curves are measured to completely characterize the amplitude crosstalk. Shishikura et al. describes a similar amplitude crosstalk measurement in [12].

The transmitter crosstalk is first measured at 3 Gb/s using a PRBS7 pattern. Fig. 16 shows a set of 13 BER curves measured with channel 5 as the reference. In this particular example, the BER curves with just one aggressor active are hardly different from the reference curve but there is a noticeable 0.5 dB penalty at BER=10^{-11} when all channels are active. The maximum crosstalk penalty for this transmitter is just under 0.6 dB and occurs on Ch. 4. Fig. 17 shows the crosstalk penalty for each channel when all twelve channels are operating simultaneously. This result is consistent with crosstalk measured on commercially available 3 Gb/s modules [31] which has been described as "under 1dB". The sum of the individual aggressor penalties does not equal the penalty for simultaneous operation. For some channels this sum exceeds the all channels penalty while for other it is less. This finding emphasizes the importance of being able to activate all 12 channels at once.

At 10 Gb/s the amplitude crosstalk penalty was significantly higher. Fig. 18 shows the reference BER curve and the simultaneous operation BER curve for Ch. 5 which is also the worst case at 10 Gb/s. The maximum crosstalk penalty is 2dB at BER=10^{-11}. Again, at 10 Gb/s we find that the sum of penalties from 11 individual aggressors does not agree with the actual penalty measured with all 12 channels operating. For comparison, Shishikura et al. reported a 0.6 dB penalty for a 4 channel transmitter with a 1 mm channel pitch [12]. Fig. 19 shows the crosstalk penalty measured on channel 5 with just a single active aggressor for all aggressor channels. This plot has the characteristic of the crosstalk being dominated by nearest neighbors except that it has an unusual feature that a next nearest neighbor, Ch. 3, contributes nearly the same penalty as a nearest neighbor, Ch. 6. This unusual behavior was correlated with the placement of the differential signal traces within the BGA pin field and the corresponding routing of the transmission lines through this pin field on the flex. In a second pass design of the flex, the transmission lines that passed through the ground cages of other channels were routed as far from the solder pads of these channels as possible. The electrical crosstalk for Ch. 7 has been measured, see Fig. 20, and found to be below -40 dB out to 20 GHz.

B. Receiver

For receiver characterization, a 2 m breakout cable is connected from the optical switch to the receiver. A single golden VCSEL is used for reference measurements while an actual transmitter module is used as the aggressor source for crosstalk measurements. The differential outputs are connected to the dual SP12T microwave switch. The non-inverting output is fed to a 10.3 Gb/s clock and data recovery module and then to a commercial error detector. The inverting output is fed to an oscilloscope.

Fig. 21 shows a representative single-ended eye diagram

from the receiver module measured at the edge of the test card. The electrical amplitude is 240 mV. The rise and fall times are 37 ps (20%-80%). The golden VCSEL was used to characterize the sensitivity of each channel. At BER=10^{-11} the sensitivity ranged from an optical modulation amplitude (OMA) of -13.0 dBm to -14.0 dBm. The plot of BER vs. OMA for all 12 channels is shown in Fig. 22. The measured sensitivity was similar for two different patterns, 2^7-1 PRBS (PRBS7) and 2^{31}-1 PRBS. However, there was a 20 ps reduction in the eye opening at BER=10^{-12} for the PRBS31 pattern due the increased baseline wander from the low cut off frequency of the receiver 's offset control, as shown in Fig. 23.

Fig. 24 shows BER vs. OMA for Ch.5 as a function of the bias voltage of the photodiode V_{bias}. There is a 0.5dB improvement in sensitivity when increasing this bias from 1.24 V to 1.40 V. The eye opening at BER=10^{-12} was also measured as a function of OMA shown in Fig. 25. Between -5 dBm and -9 dBm there is an 18 ps reduction in the timing margin. At a constant OMA = -5.0 dBm the eye opening varied from 45 ps to 55 ps for all 12 channels (PRBS7). For comparison, the test equipment used to drive the golden VCSEL has an eye opening of 60 ps.

The amplitude crosstalk was measured using a golden VCSEL to characterize the channel under test and the transmitter module was used for the aggressor channels. The use of a golden VCSEL allows the measurement to be independent of the crosstalk penalty within the transmitter. The measured crosstalk penalty with all channels operating varied from 2.5 dB to 3.5 dB across the 12 channels. Shishikura et al. report a 4 dB receiver crosstalk penalty for a 4 channel receiver measured in a similar manner [12]. Upon examining the crosstalk penalty due to individual aggressors it was noted that some channels could generate penalties that were larger or comparable with that from a neighboring channel. It was observed that the measured crosstalk was a a function of the optical amplitude of the signal input to the aggressor channel. If a particular channel is in between channels with very large OMA then the resulting penalty for that channel will be higher than for a channel with low OMA neighbors. Fig. 26 shows the crosstalk penalty measured on Ch. 2 as a function of the OMA of the neighboring channels 3 through 12. In this example, the crosstalk from Ch. 3 is dominant, with neighboring channels not contributing to the penalty until their OMA value exceeds -5 dBm.

C. *Link Results*

Having separately characterized the transmitter and receiver modules, a link was put together and operated with all channels running simultaneously. The link consisted of a 10 m jumper of standard 12 Ch. 50/125 MMF ribbon fiber connected to a 1m jumper, a 300 m of OM-3 MMF, another 1m jumper, a 2 m breakout cable, the fiber optic switch, and finally a second 2 m breakout cable from the switch to the receiver. The total link length was approximately 316 m and contained 3 MTP® to MTP® connectors. The total link power dissipation is 2.9 W. A clock and data recovery circuit is used to recover a full rate clock from each channel for BER testing.

The crosstalk in this link configuration was remeasured since the optical transmitter is the real transmitter module which has more jitter and internal crosstalk that is not present in the golden VCSEL used to characterize the receiver. In this case the crosstalk penalty varied from 1.5 dB to 4.0 dB and was correlated with the OMA values on neighboring channels.

A study was performed to see if lowering the transmitter's OMA could improve link margin by lowering the total link crosstalk penalty. Table 1 shows the change in link margin on channels 2 and 6 for three different transmitter OMA settings. In general it was found that the increase in OMA exceeded the increase in crosstalk penalty but with diminishing returns. For example, increasing the transmit optical amplitude by 2 dB also increased the link crosstalk by 0.7 dB, resulting in a net 1.3 dB gain in link margin.

In the 316 m link configuration, all channels were observed to run error free for time periods observed under 1 hour. Two channels were selected for long term observation. One of these accumulated 9 errors overnight (18 hrs) for a BER = 10^{-14}. The other accumulated 9 errors in 68 hours (the weekend) for a BER=$4x10^{-15}$.

VI. Conclusion

A custom low cost tester was developed to help lower the cost of development of parallel optical interconnects. Using this tester, a 120Gb/s optical link has been demonstrated and characterized under simultaneous operation for crosstalk and link margin. The BER was observed to be less than 10^{-12} for all 12 channels under simultaneous operation over 316m. This meets the BER requirements for optical interconnects used in node to node computing connections. These results show that the existing 30 Gb/s optical interconnect packaging can be extended to 120 Gb/s and possibly up to 180 Gb/s although the scaling of crosstalk with data rate needs to be investigated.

Acknowledgment

The authors wish to acknowledge John Crow, Jeff Kash, Modest Oprysko and Stan Swirhun for the initiation and supoort of this project. Additionally, the authors are grateful for the construction support for the testing tower from the IBM T. J. Watson Research Center CSS staff.

References

[1] www.infinibandta.org
[2] J.D. Crow, Crowsquill Consulting, "Optical Interconnects Today," Proceedings of 14th Annual Workshop on Interconnections within High Speed Digital Systems, May 4-7, 2003, Santa Fe NM, paper PLE2.
[3] R. Michalzik, K.J. Ebeling, M. Kicherer, F. Mederer, R. King, H. Unold, and R. Jager, "High-performance VCSELs for optical data links", IEICE Transactions on Communications v.E85-B n. 5 May 2001 p.1255-1264.
[4] Honeywell 10Gb/s white paper **http://content.honeywell.com/vcsel/pdf/10GWhitePaper.pdf**

[5] IEEE Standard 802.3ae™, IEEE (Piscataway, New Jersey, June 2002).

[6] See for example the websites of Infineon Inc., Picolight Inc, and JDSU Inc.

[7] C. Cook, J. E. Cunningham, Al Hargrove, G. G. Ger, K. W. Goossen, W. Y. Jan, H. H. Kim, R. Krause, M. Manges, M. Morrissey, M. Perinpanayagam, A. Persaud, G. J. Shevchuk, V. Sinyansky, and A. V. Krishnamoorthy, "A 36-channel parallel optical interconnect module based on optoelectronics-on-VLSI technology", IEEE Journal on Selected Topics in Quantum Electronics, Volume: 9 , Issue: 2 , March-April 2003 p. 387 - 399

[8] See for example the websites of Infineon Inc., Picolight Inc, and Agilent Inc.

[9] Giboney, K.; Simon, J.; Mirkarimi, L.; Law, B.; Flower, G.; Corzine, S.; Leary, M.; Tandon, A.; Kocot, C.; Rana, S.; Grot, A.; Kun-Jing Lee; Buckman, L.; Dolty, D., "Next-generation parallel-optical data links", The 14th Annual Meeting of the IEEE Lasers and Electro-Optics Society, 2001. LEOS 2001.,Volume: 2 , 12-13 Nov. 2001 p. 859 - 860

[10] A. Schild, H.-M. Rein, J. Müllrich, L. Altenhain, J. Blank, and K. Schrödinger, "High-gain SiGe transimpedance amplifier array for a 12 x 10 Gb/s parallel optical-fiber link," IEEE J. Solid-State Circuits, vol. 38, pp. 4-12, Jan. 2003.

[11] Schild, A. Rein, H.-M. Mullrich, J. Altenhain, L. Blank, J. Schrodinger, K., IEEE Radio Frequency Integrated Circuits Symposium, RFIC, Digest of Technical Papers 2002 p.89-92

[12] Shishikura, M.; Ban, T.; Ichikawa, H.; Ido, T.; Takahashi, M.; Nakahara, K.; Nomoto, E.; Matsmoka, Y.; Ishikawa, K.; Ito, K.; Takeyari, R.; Sano, H.; Hirose, T.; Takahashi, H.; Nagara, T.; Chiba, H.; Irikura, S.; "10 Gbit/s per channel parallel optical transmitter and receiver modules for high-capacity interconnects", Proceedings. 53rd Electronic Components and Technology Conference, 2003., May 27-30, 2003 p. 1040 - 1045

[13] Sakamoto, T.; Sato, N.; Koike, S.; Hadama, K.; Kukutsu, N.; Hashimoto, T.; Ohyama, T., "4 Channel x 10 Gbit/s parallel optical module for short reach optical links", LEOS 2003. The 16th Annual Meeting of the IEEE Lasers and Electro-Optics Society, 2003. , Volume: 2 , 2003 p. 567 - 568

[14] K. Kurata, I. Hatakeyama, K. Miyoshi, T. Shimuizu, J. Sasaki, M. Kurihara, K. Yamamoto, "Opto-electronic packaging techniques for interconnection", LEOS 2003. The 16th Annual Meeting of the IEEE Lasers and Electro-Optics Society, 2003. ,Volume: 1 , 26-30 Oct. 2003 p.364 - 365

[15] Harting Inc. website on Pareos. http://www.harting.com

[16] Michalzik, R. Giaretta, G. Goossen, K.W. Walker, J.A. Nuss, M.C., "40 Gb/s coarse WDM data transmission with 825 nm wavelength VCSELs over 310 m of high performance multimode fiber", ECOC 2000. 26th European Conference on Optical Communication 2000 P33-4 vol.4

[17] See for example the websites of Xanoptix Inc., Aralight Inc, and Terraconnect Inc.

[18] Cunningham, J.E.; Krishnamoorthy, A.V.; Jan, W.Y., "36-channel, 120Gigabit/s parallel optical interconnect", LEOS 2003. The 16th Annual Meeting of the IEEE Lasers and Electro-Optics Society, 2003. , Volume: 1 , 26-30 Oct. 2003 p.156 - 157

[19] A. Joseph et al, "A 0.18µm BiCMOS Technology Featuring 120/100GHz (fT/fmax) HBT and ASIC-Compatible CMOS Using Copper Interconnect", IEEE Proc. BCTM, pp. 143-146, 2001

[20] E.M. Cherry and D.E. Hooper, "The design of wideband transistor feedback amplifiers," Proc. IEE, vol. 110, pp. 375–389, Feb. 1963.

[21] TIA/EIA-644 and IEEE 1596.3 Standards.

[22] W. Klein and B.-U. H. Klepser, "75 GHz bipolar production technology for the 21st century," in Proc. ESSDERC, 1999, pp. 88-94.

[23] http://www.snapoptics.org

[24] MEG-Array® Connector System, Datasheet and Application Notes, FCI USA (Etters, Pennsylvania, 2002). http://www.fciconnect.com

[25] Pyralux R AP, Technical Information, DuPont (Wilmington, Delaware, 2003). http://www.dupont.com/fcm/products/H-73241.html

[26] C. Schuster, C.W. Baks, D.M. Kuchta, Y.H. Kwark, L. Graham, "Electrical Interconnect Design and Optimization for 120 Gbps Parallel Optical Transmitter Module and Test Station", Proceedings of the 7th IEEE WORKSHOP ON SIGNAL PROPAGATION ON INTERCONNECTS.

[27] L. Shan, D. Kuchta, Y. Kwark, "A method of multi-node thermal network extraction for SPICE compatible simulations", Proceedings. 53rd Electronic Components and Technology Conference, 2003., May 27-30, 2003 p. 1631 - 1635

[28] Dolti, D.W., "Multi-channel optical interconnects for short-reach applications", Proceedings. 53rd Electronic Components and Technology Conference, 2003., May 27-30, 2003 p. 1032 - 1039

[29] TIA/EIA Fiber Optic Test Procedure 203, "Collection and Reduction of Two-dimensional Nearfield Data."

[30] TIA-492AAAC: "Detail Specification for 850-nm Laser-Optimized, 50-µm Core Diameter/125-µm Cladding Diameter Class Ia Graded-Index Multimode Optical Fibers"

[31] Private Communication with several 3Gb/s vendors

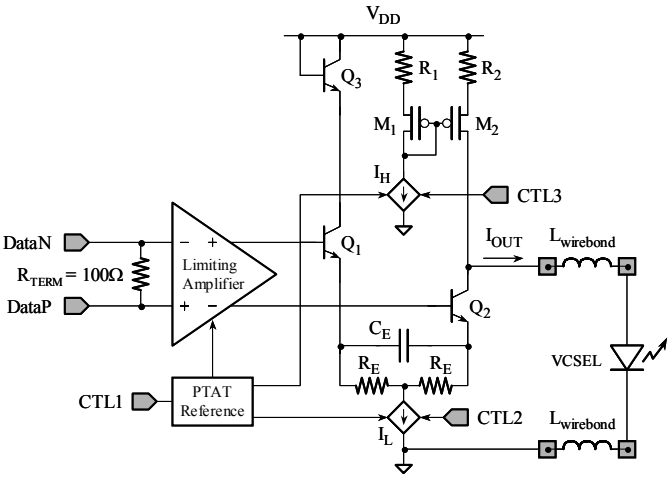

Fig. 1 Block diagram of an individual channel of VCSEL driver IC.

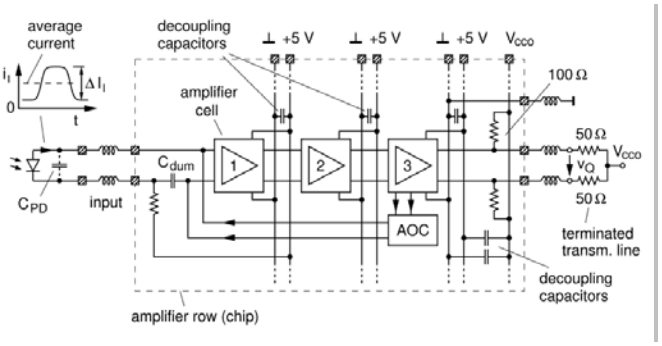

Fig. 2 Block diagram of an individual channel of the receiver IC

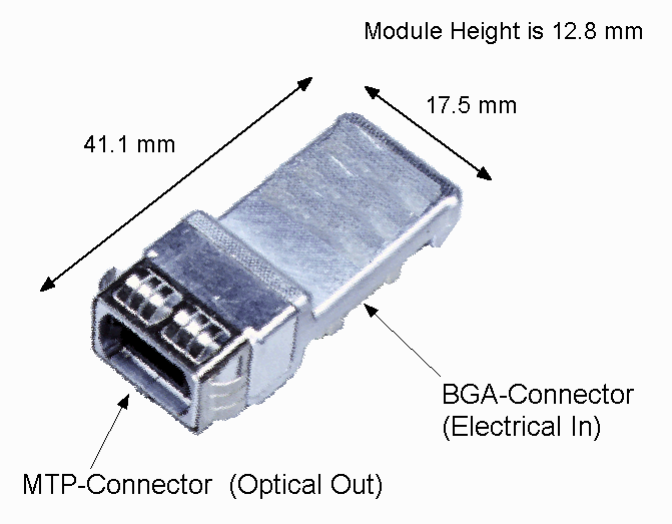

Fig. 3 Picture of fully assembled optical module, without heatsink, for both transmitter and receiver.

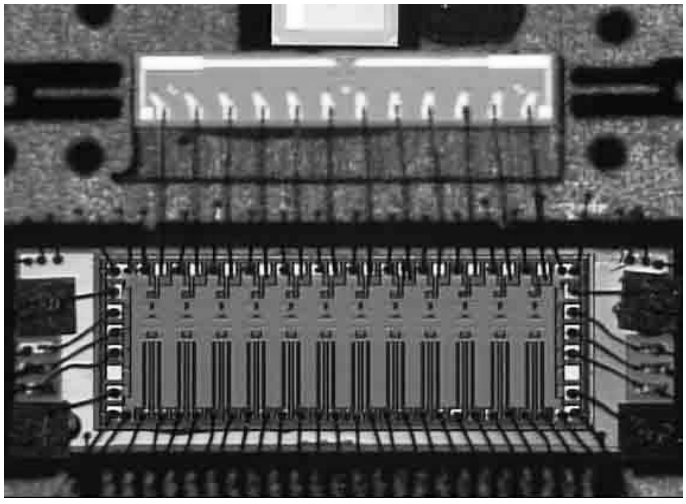

Fig. 4 Transmitter internal assembly showing wirebonds and placement of driver (foreground) and VCSEL array.

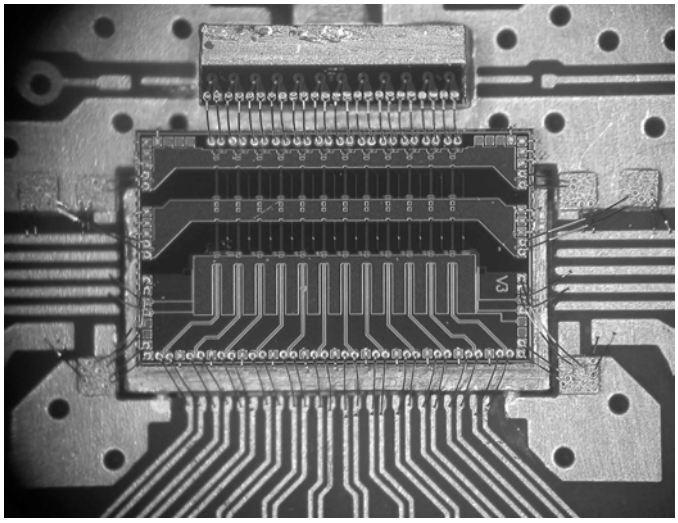

Fig. 5 Photo of assembled receiver IC and photodiode array on flex circuit showing wirebonds and decoupling capacitors.

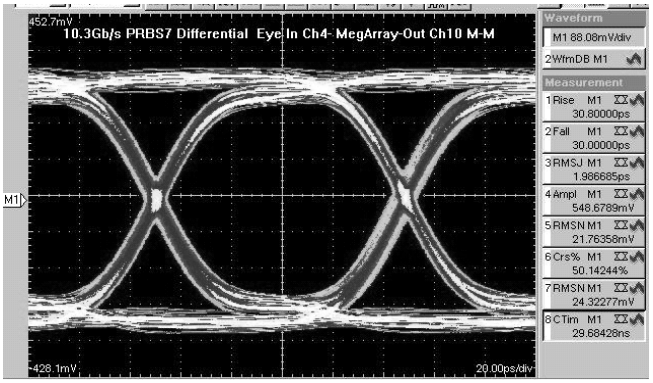

Fig. 6 Eye diagram at 10.3 Gb/s through MEG-Array connector sandwiched between two test cards. The voltage and time scales are 88 mV/div and 20 ps/div.

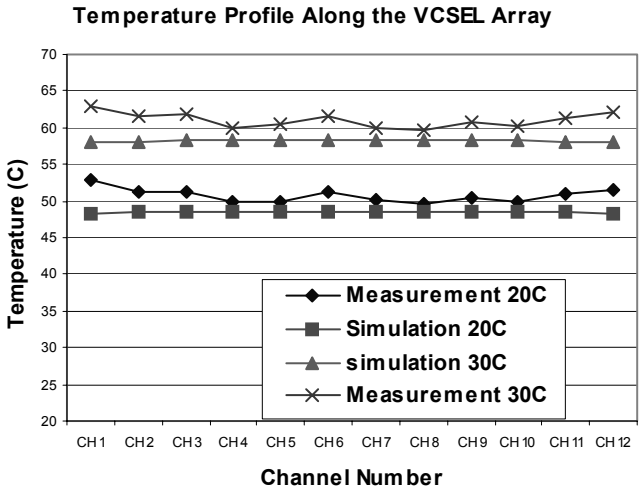

Fig. 7 Temperature variation along the VCSEL array and comparison with a 3D simulation for two different ambient temperatures 20° C and 30° C.

Fig. 8 Photograph of the custom 120 Gb/s testing tower. The optical transmitter is visible in the center and has a 12 fiber ribbon attached.

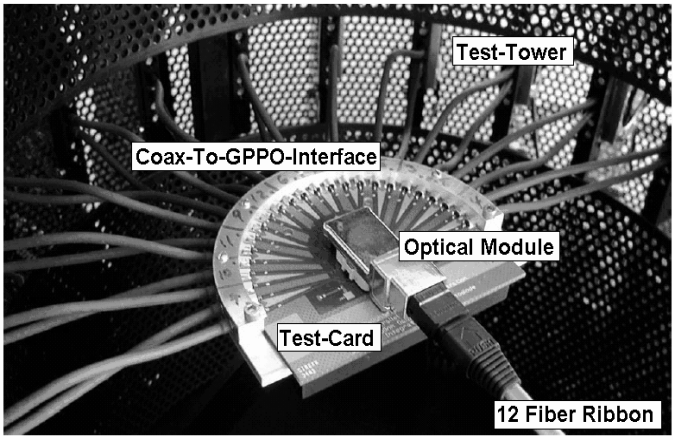

Fig. 9 Testing tower test card which interfaces between the tester and the device under test.

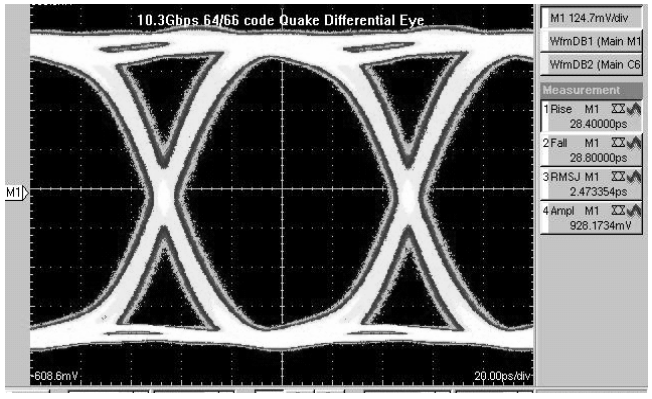

Fig. 10 Differential electrical eye generated at the 10 Gb/s outputs of the tester. The rise and fall times are < 30 ps and the jitter is < 2.5 ps (rms). The voltage and time scales are 125 mV/div and 20 ps/div.

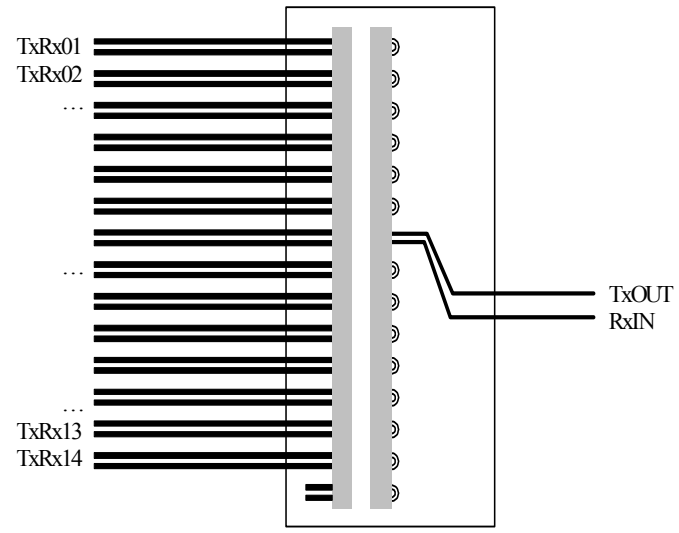

Fig. 11 Schematic drawing of custom optical switch inserted into a link to allow one channel to be individually brought out to test equipment such as an optical attenuator while all other channels pass through from input to output.

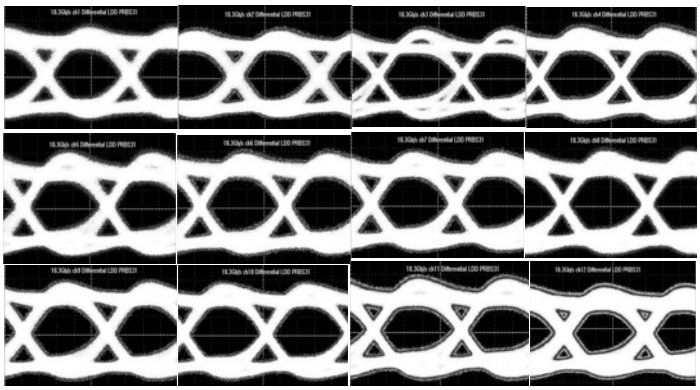

Fig. 12 Optical eye diagrams of transmitter at 10.3125 Gb/s with a PRBS31 pattern before the optical coupler is attached.

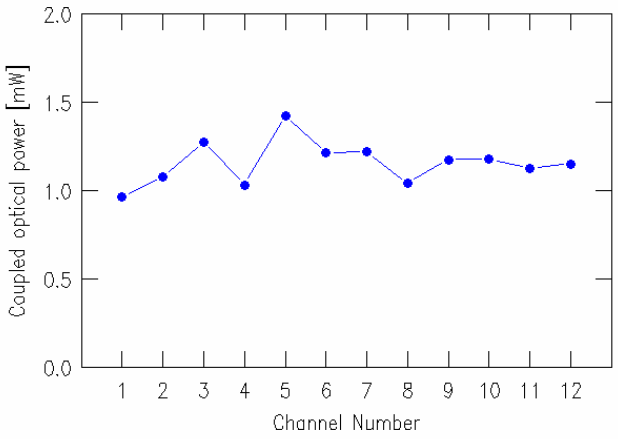

Fig. 13 Average optical power of a completed transmitter coupled into 50/125 MMF across 12 channels showing a variation of +/-15% (1.5dB).

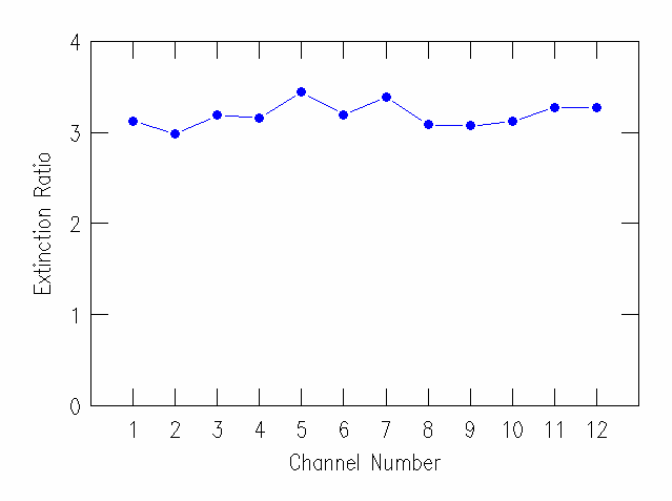

Fig. 14 Extinction ratio across 12 channels showing a variation of +/-10% (1dB).

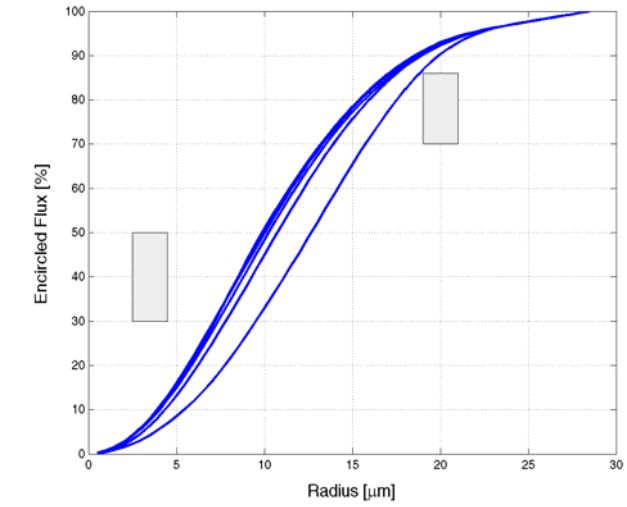

Fig. 15 Encircled flux measurement of all 12 transmitter channels.

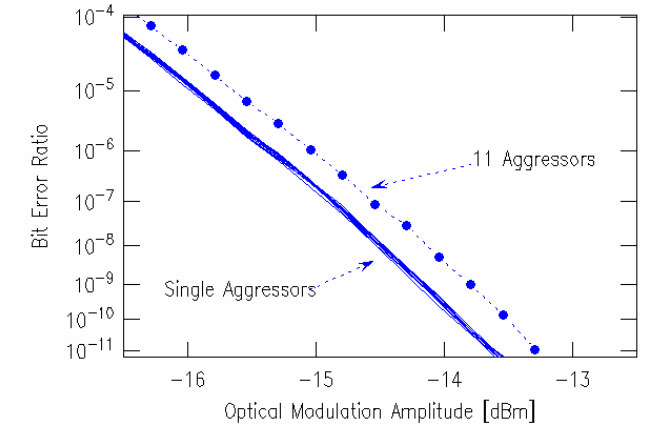

Fig. 16 Transmitter crosstalk BER curves for Ch. 5 at 3 Gb/s. The penalty for simultaneous operation is < 0.5 dB.

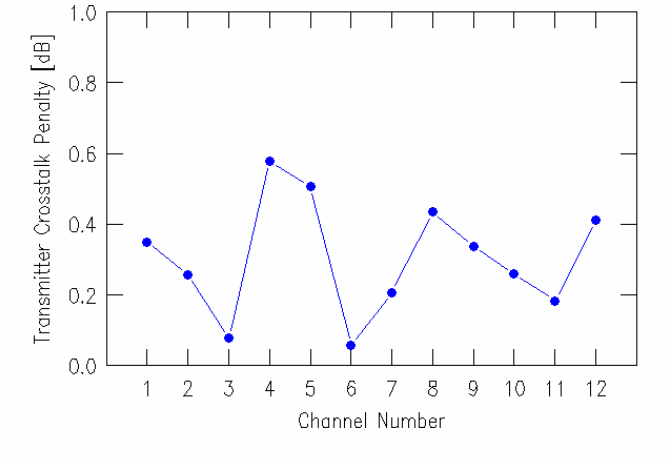

Fig. 17 Transmitter crosstalk penalty by channel for simultaneous operation at 3 Gb/s.

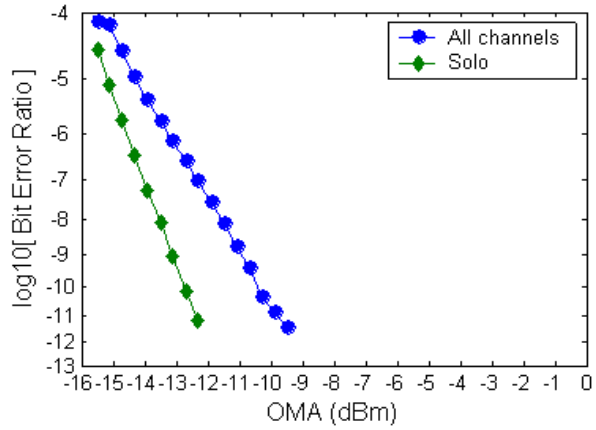

Fig. 18 Transmitter crosstalk BER curves for Ch. 5 measured at 10Gb/s, solo and with all channels operating.

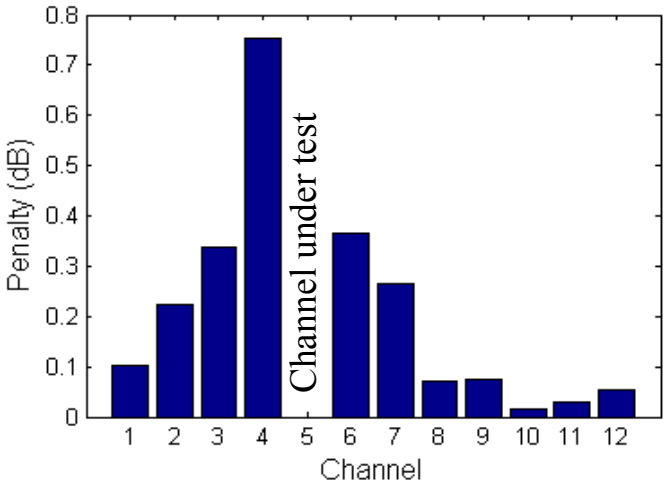

Fig. 19 Transmitter crosstalk penalty by neighboring channel for channel 5 measured at 10 Gb/s.

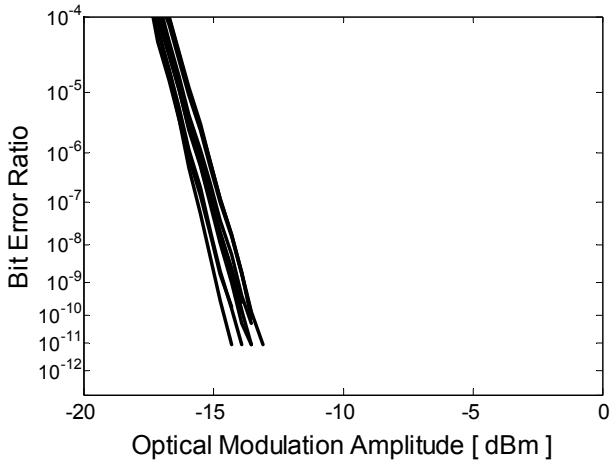

Fig. 22 BER vs. OMA for all 12 receiver channels measured in solo operation.

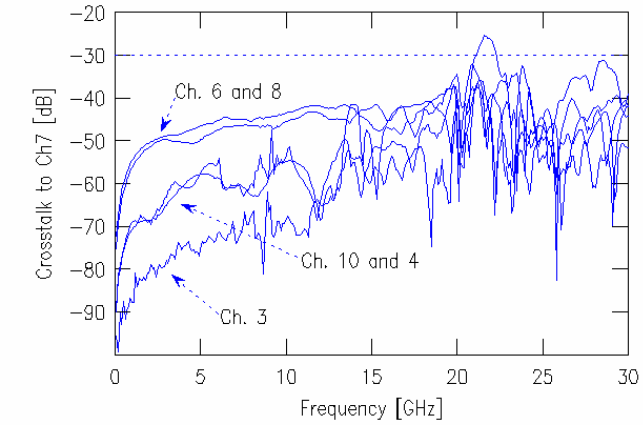

Fig. 20 Electrical crosstalk measured on Ch. 7 through the test card, MEG-Array and redesigned flex using a 4-port vector network analyzer. The top two traces are crosstalk from nearest neighbors. Channels 10 and 4 are nearest neighbors in the BGA connector.

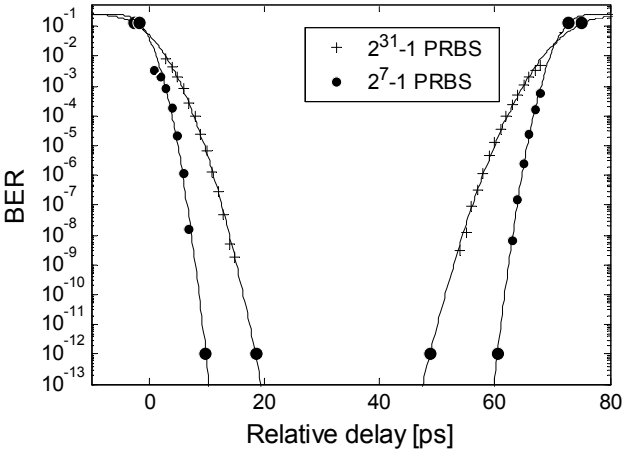

Fig. 23 Receiver eye opening (bathtub curve) at the receiver output in solo operation for PRBS7 and PRBS31 at -5.0 dBm OMA. There is a 20 ps reduction in the eye opening for the PRBS31 pattern.

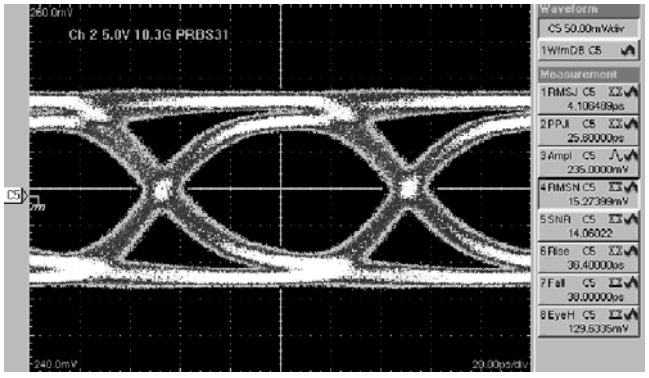

Fig. 21 Representative eye diagram at the output of the receiver. The voltage and time scales are 50 mV/div and 20 ps/div.

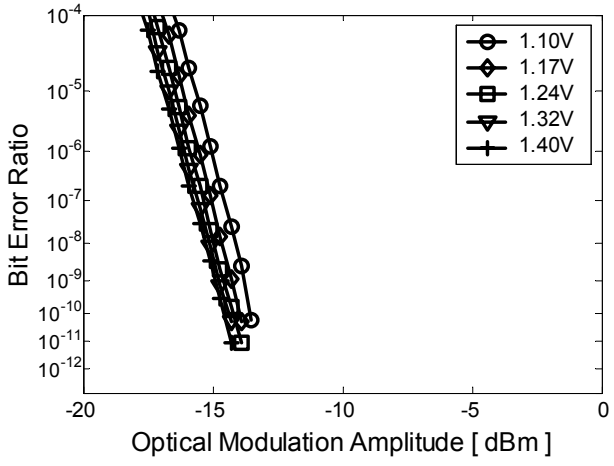

Fig. 24 BER vs. OMA for a representative channel of the 12 channel receiver in solo operation at different bias voltages, V_{bias}, of the photodiode.

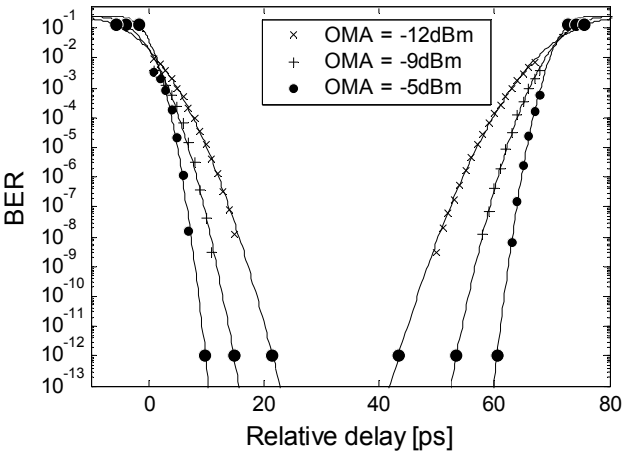

Fig. 25 Receiver eye opening in solo operation for different OMA values of -5, -9, and -12 dBm. The pattern is PRBS7.

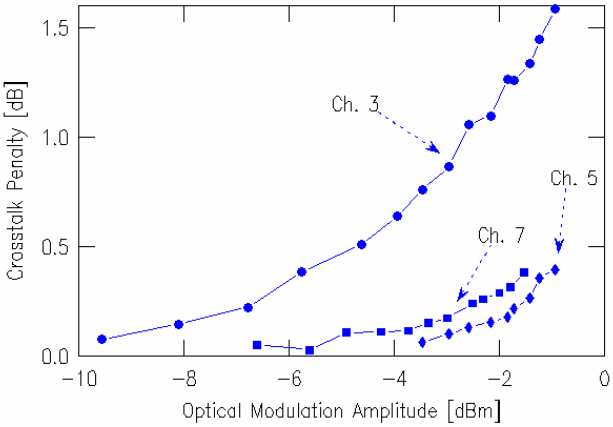

Fig. 26 Receiver crosstalk penalty for channel 2 as a function of aggressor channel OMA for selected channels.

Ave. OMA change [dB]	Δ penalty Ch6 [dB]	Δ gain Ch6 [dB]	Δ penalty Ch2 [dB]	Δ gain Ch2 [dB]
2.1	0.8	1.3	0.8	1.3
1.3	0.9	0.4	1.0	0.3
3.4	1.6	1.8	1.8	1.6

Table 1 Change in link margin with change in transmitter OMA.

Daniel M. Kuchta (S'85-M'92-SM'97) is a Research Staff Member in the Communication Technology Department at the IBM Thomas J. Watson Research Center. He received B.S., M.S., and Ph.D. degrees in electrical engineering and computer science from the University of California at Berkeley in 1986, 1988, and 1992, respectively. He subsequently joined IBM at the Thomas J. Watson Research Center, where he has worked on high-speed VCSEL characterization, multimode fiber links, and parallel fiber optic link research. Dr. Kuchta is an author or coauthor of nine patents and more than 25 technical papers.

Young H. Kwark Biography not available at this time.

Christian Schuster Christian Schuster (S'98 - M'00) received the Diploma degree in physics from the University of Konstanz, Germany, in 1996, and the Ph. D. degree in electrical engineering from the Swiss Federal Institute of Technology (ETH), Zurich, Switzerland, in 2000. From 2000 to 2001 he was with ISE Integrated Systems Engineering AG, Zurich, Switzerland, where he was responsible for development and support of an electromagnetic field solver software. Currently he is a Research Staff Member at the IBM T. J. Watson Research Center, Yorktown Heights, NY, where he is involved in high-speed optoelectronic package and backplane interconnect modeling, design, and optimization. Dr. Schuster is member of the German Physical Society (DPG). He received the IEEE Transactions on EMC Best Paper Award in 2001.

Christian Baks Biography not available at this time.

Chuck Haymes received his undergraduate degree in electrical engineering from Northwestern University and his masters degree in computer engineering from North Carolina State University. He is an Advisory Engineer at the IBM TJ Watson Research Center. He has been granted five US patents, one Japanese patent and one European Union patent, all of which relate to networking and digital communication systems.

Jeremy Schaub Biography not available at this time.

Petar Pepeljugoski is a Research Staff Member in the Communication Technology Department at the IBM Thomas J. Watson Research Center. He received his B.S. from Ss. Cyril and Methodius University in Skopje, Macedonia in 1982, his M.S. from University of Belgrade, Yugoslavia in 1986 and his Ph.D. from the University of California at Berkeley in 1993. He joined IBM in 1994, where his research work included design, modeling, prototyping and characterization of high speed multimode fiber LAN links and parallel interconnects. His modeling tools were used by the Telecommunication Industry Association in the development of the next generation fiber. He was involved in the development of the 10 Gigabit Ethernet Standard (IEEE 802.3ae) from its beginning and has been awarded IEEE recognition by the 10 Gigabit Ethernet Alliance for his contributions. Dr. Pepeljugoski was also awarded Certificate of Appreciation by the Telecommunication Industry Association for his contribution to the Working Group on Modal Bandwidth of Multimode Fibers and the development of U.S. standards for Fiber Optic Technology.

Lei Shan Biography not available at this time.

Richard John Biography not available at this time.

Dennis Rogers Biography not available at this time.

Mark Ritter Dr. Ritter received a B.S. degree in physics from Montana State University in 1981 and M.S., M.Phil. and Ph.D. degrees from Yale University in 1987. In the fall of 1986, Dr. Ritter took a post-doctoral research position in the physical sciences department of the IBM T.J. Watson Research Center in Yorktown Heights, New York. His post-doctoral work with Dr. David Awschalom centered on nonlinear optical properties of fluid-filled porous media and carrier dynamics in diluted magnetic super lattice structures. Dr. Ritter was permanently hired as a research staff member in January of 1989. Since 1993 he has been manager of a group designing analog and digital circuits for data communications. His group has contributed to the Fibre Channel standard for storage-area networks and to the 802.3ae 10Gb/s Ethernet standard, as well as to the IR wireless (IrDA) standard. The group work has also resulted in a number of IBM products, including Fibre Channel transceivers, IrDA transceivers, and, most recently, 10 Gb/s Ethernet transceivers where his group contributed to the design of the analog front-end circuits at 10Gb/s. He and his group are exploring silicon-germanium circuits for communications links at 40 Gb/s and beyond, and are exploring silicon circuits and technologies suitable for future Terabit communications links within switches and servers. Dr. Ritter was the recipient of the 1982 American Physical Society Apker Award, three IBM Outstanding Innovation awards, and several Research Division and Technical Group Awards. He is an author or coauthor of 30 technical publications and a co-inventor of 14 U.S. patents.

Jack Jewell Biography not available at this time.

Luke Graham Biography not available at this time.

Daniel Kucharski (S '00) received the B.S. degree in electrical engineering from RIT in 1998 and joined IBM Microelectronics as a Development

Engineer, where he worked on high-speed CMOS I/O designs. In 2000 he joined Broadband Communications Research Lab at Cornell University, where he is currently pursuing a Ph.D. degree. He spent the summers of 2001–2003 at IBM's T.J. Watson Research Center, where he designed laser driver arrays in SiGe and advanced CMOS. His research interests focus on high-speed low-power optical communication circuits.

Karl Schrodinger Karl Schrödinger was born in Straubing, Germany, in 1953. He received his Diploma in Electrical Engineering from the Technical University Munich in 1980. He joined Siemens in 1980, starting a 200 MBit/s link development in the former Data Systems Group. Later he joined the Siemens Networking Unit and developed electronic circuits, integrated circuits, optical modules with data rates between 200MBit/s and 10 GBit/s within the Fiber Optics Group of Siemens AG and later Infineon Technologies AG. He holds several positions such as manager for electronics, project manager for Gigabit Ethernet Modules and manager for integrated circuit development. Currently he is responsible for high speed IC development within the Fiber Optics Group in Berlin/Germany.

Alexander Schild Alexander Schild (S '97) was born in Bonn, Germany, in 1971. He received the Dipl.-Ing. degree in electrical engineering from Ruhr-University Bochum, Germany, in 1999. Then he joined the research group Semiconductor Devices (Arbeitsgruppe Halbleiterbauelemente) at Ruhr-University Bochum as a Research Assistant in 1999. His main field of research is the design of high-speed bipolar circuits for fiber optic transmission systems.

Hans-Martin Rein Hans-Martin Rein (M '92 - SM '99 - F '01) received the Dipl.-Ing. and Dr.-Ing. degrees in electrical engineering from the University of Stuttgart, Germany, in 1964 and 1968, respectively. From 1968 to 1973, he was with the Semiconductor Division of AEG-Telefunken in Heilbronn, Germany. As head of a department he was responsible for the development of high-speed bipolar circuit families and of advanced MOS/CMOS technologies. Since 1973 he has been Professor at Ruhr-University Bochum, Germany, where he is teaching integrated circuits. His main fields of research are the design and optimization of high-speed analog and digital Si-bipolar ICs (especially for optical communications and microwave applications) with operating speeds up to 60 GHz and 60 Gb/s, respectively, as well as modeling of very fast bipolar transistors (incl. SiGe HBTs) and of on-chip parasitics (e. g. substrate coupling). He is author and co-author of more than 150 technical papers and of a textbook on integrated bipolar circuits.