AdvancedTCA Backplane Tester

Alexandra Dana Oltean, Brian Martin

"POLITEHNICA" University Bucharest Romania

and

CERN 1211Geneva 23 Switzerland

Email: alexandra.oltean@cern.ch, brian.martin@cern.ch

***Abstract*— The "Advanced Telecom Computing Architecture" (AdvancedTCA) is a modular standard chassis based system designed to support the needs of carrier class telecommunication applications. It is defined by a set of industry standards under the direction of the PICMG group. One early deployment of the standard technology has been a 10 Gigabit Ethernet switch developed in the framework of the EU funded ESTA project.**

In order to study the practical aspects of high speed Ethernet switching at 10 Gigabit and above and to validate the signal integrity of the AdvancedTCA backplane, we developed a Backplane Tester. This system is able to run pseudo-random bit sequence (PRBS) traffic at 3.125 Gbps over every link on the AdvancedTCA backplane simultaneously, and to monitor any possible connectivity failure immediately in terms of the link and slot positions inside the chassis.

In this paper, we describe the design and the practical architectural hardware and software aspects of the AdvancedTCA Backplane Tester. We also present the results of Bit-Error Rate (BER) test measurements performed with our Tester on a fully populated 10 Gigabit traffic AdvancedTCA Backplane.

I. Introduction

SINCE the ratification of the Advanced Telecom Computing Architecture (AdvancedTCA®) standard in 2003 [1], many high-availability applications in the telecom industry were selecting this architecture for their designs. A 10 Gigabit Ethernet switch, built within the ESTA EU project [2], based its implementation on the AdvancedTCA backplane, on which multiple 10 Gigabit (Gbps) Ethernet Line cards and two Switch Fabric cards were plugged.

The AdvancedTCA backplane was subject to extensive and professional simulation, testing and demonstration before meeting the specifications included in the standard. However the actual construction of the backplane can be itself subject to printed circuit board (PCB) manufacturing control problems, which could greatly affect its quality control. For the 10 Gigabit Ethernet switch, as for other commercially produced equipments, a low quality backplane combined with marginal silicon performance could cause low level errors in the system. Tracking down the cause of these errors and fixing them in the final product, can sometimes take as long as the original design cycle.

It would greatly simply the commissioning process if a tool were available that could validate and characterize the backplane sufficiently to demonstrate that it could not be the source of any observed low level errors.

We first considered commercially available backplane testers but these were only connectivity testers operating at DC or very low frequencies and would not detect high speed signal integrity problems. There are of course programmable high speed data generators and checkers available but they have a very high level of functionality that is not necessary for this application, and a correspondingly high cost per channel. It was therefore decided to develop our own tester which would serve not just the requirements for prototype testing, but also as a means to check any AdvancedTCA production backplane samples.

The tester was designed as a high-speed data traffic generator at 10 Gbps, which specifically addresses and tests the most critical aspects of the AdvancedTCA backplane. In the following paragraphs, the design and the practical architectural hardware and software aspects of the AdvancedTCA Backplane Tester will be described. The results of Bit-Error Rate (BER) Test measurements performed with the Tester at design speed on a fully populated and fully-loaded AdvancedTCA Backplane will

also be presented.

II. Backplane Tester Hardware Architecture

A. Traffic Generation in the Backplane Tester

A Dual-Star Fabric Interface topology, offering a redundant backplane fabric scheme, has been selected for the transmission of the high-speed data stream across the AdvancedTCA backplane. This topology (presented in Fig. 1) requires two dedicated slots for the Switching Cards (Hub1 and Hub2 boards) while twelve other slots in the AdvancedTCA chassis are dedicated to the Line cards (Node boards). The redundant connection between each Node and both Hub boards is achieved via two Fabric Channels, with a raw throughput of 12.5 Gbps each. The Fabric Channel itself is based on four differential pairs per direction. In the 10 Gigabit Ethernet switch, each Fabric Channel has been enlarged beyond the requirement of the AdvancedTCA standard. This extension was based on two extra differential pairs per direction so that the effective throughput of each Fabric Channel became 15 Gbps. Thus, it results a total of 25 Fabric Channels, which occupy 300 differential pairs, each operating at 3.125 Gbps on the AdvancedTCA backplane.

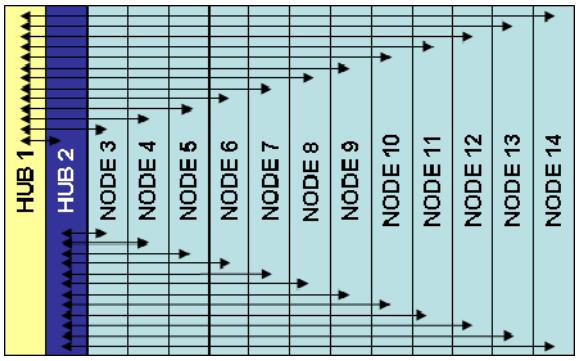

Fig. 1. Fabric Interface Dual Star Topology

In order to provide traffic for saturating all the available pairs on the backplane, the Alaska 88X2040, a 10 Gigabit Ethernet device from Marvell, has been selected [3]. This device uses the same transmit/receive pad technology as the switching fabrics in the actual 10 Gigabit Ethernet switch and has built-in BER circuitry which permits to generate and check traffic without external devices. The traffic can be provided using the Alaska 88X2040 to generate Jitter test patterns or Pseudo-Random Bit Sequence (PRBS) [4] for stressing various aspects of system or component performance in the backplane application.

The generated data is sent to the differential pairs on the backplane and on the end of each pair, the received data is de-serialized, aligned, de-skewed, decoded and finally checked. The BER feature of the Alaska device permits allows for the independent testing of each differential pair connection by monitoring the errors appearing on each pair involved in the backplane test. The implementation of Alaska devices in the AdvancedTCA backplane tester system is reported in Fig. 2. In Fig.2, we show only the star connection starting from Hub1 for clarity, but an identical (redundant)

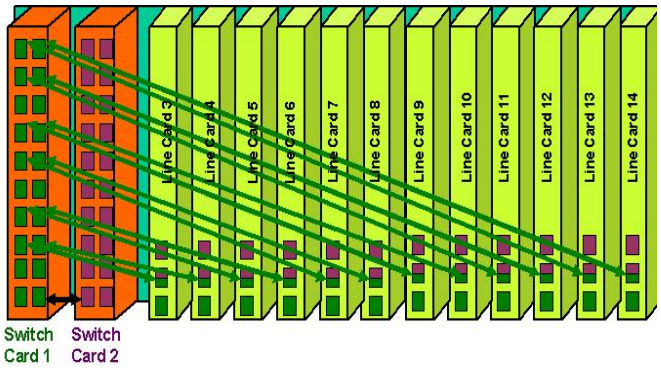

Fig. 2. Marvell Alaska in the Backplane Tester

star connection is made to the Hub2 board.

Since each Alaska device embeds four differential drivers and four differential receivers operating at 3.125 Gbps each, three of them are required to saturate the two redundant Fabric Channels of each Node board. To achieve connection from a Hub board towards all the slots in the chassis, twenty Alaska devices are used for each of the Hubs. Thus, in the full backplane tester system there are 76 Alaska devices altogether with twenty for each Hub board and three for each Node board.

B. AdvancedTCA Backplane Tester Control

For controlling the traffic generation over the AdvancedTCA backplane we used the Base Interface, defined in the standard with a dual-star topology as well. In order to reduce board design, costs and testing procedures, we decided for a single point of control for the whole tester system, implemented in Hub1. The Base Interface Channels issued from Hub 1 are then used to distribute control signals to all Node boards and to Hub2.

We decided to use the Lantronix (DST-LX-001) microcontroller [5] as control device for the AdvancedTCA

backplane tester system. This device offers much embedded functionality at a relatively low cost, such as a Real-Time Operating System, TCP/IP stack, Ethernet interface etc., and a reference development platform is also available. The controller is operating with 256KByte on-chip and 1MByte on-board RAM, it is associated with on-board parallel and serial flash memories.

For generating traffic and maintaining error statistics over each differential pair on the backplane, the Marvell's devices are programmed and controlled via the MDIO interface [4]. For their control over MDIO, three lines are to be used: a MDIO bi-directional data line (formed by the output 'MDO' and the input 'MDI'), a Management Data Clock ('CLOCK') and a reset line ('RESET'). Each differential pair on the backplane has controlled impedance tracks designed for carrying 10/100/1000 Mbps Ethernet signalling so control signal integrity is not a problem. However, the control port lines need to be buffered to differentially drive the Base Interface Channels. Since the MDIO lines are bidirectional, the direction to which the data line buffers are driving the signal is specified by a 'RD' signal provided by the microcontroller. It is then necessary, in addition to the three control lines, to send this fourth signal pair over the backplane as well.

There was a further complication in that the MDIO frame could only support five bits of addressing for the physical devices ports. That made a maximum of 32 addressable Marvell devices that can be located on an individual MDIO bus. Since there are twenty MDIO ports on each Hub and three MDIO ports on each Node board, the control point needs to manage more than one single MDIO chain. We partitioned the design so that four MDIO busses are used: two of the busses for controlling each Hub and the other two for the rest of the twelve Node boards.

The control architecture (Fig. 3) has been defined so that a single PCB can be used either for the Nodes or for the Hubs design, by populating the boards in a different way.

The Hub boards are fully populated with twenty Alaska 88X2040 devices (Serdes1 to Serdes20), whereas Node boards are partially populated with only three of them (Serdes1 to Serdes3) to handle two Fabric Channels. The control signals MDI, MDO, CLOCK, RD and RESET are sent from the Lantronix controller on Hub1, and they are received on Base Channel 1 of Node boards and on Base Channel 2 of Hub 2.

An array of jumpers is used to define whether the board is a Hub or a Node and thus whether to drive or to receive the various control and data lines.

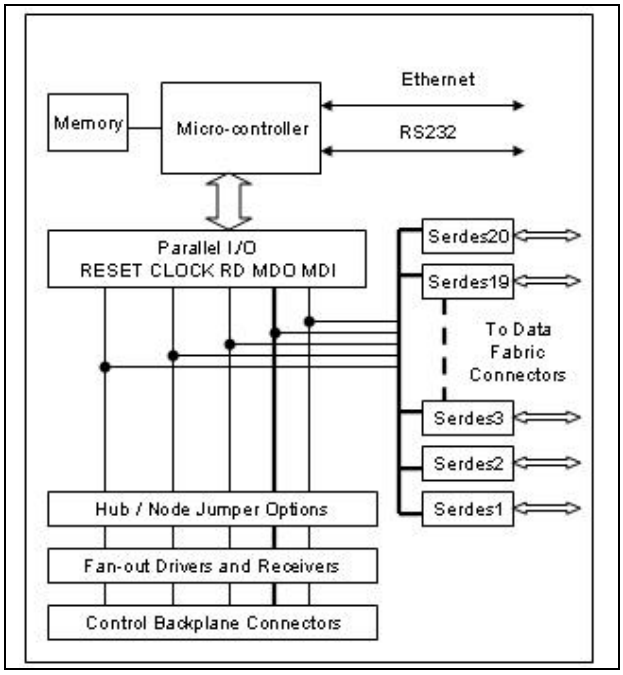

Fig. 3. AdvancedTCA Backplane Tester Control Architecture

The AdvancedTCA backplane tester (previously described) has been built and the Fig. 4 shows one board populated for Hub 1 purpose. Dimensions are 322.5 mm x 280 mm. A 14-layers stack-up was used for production.

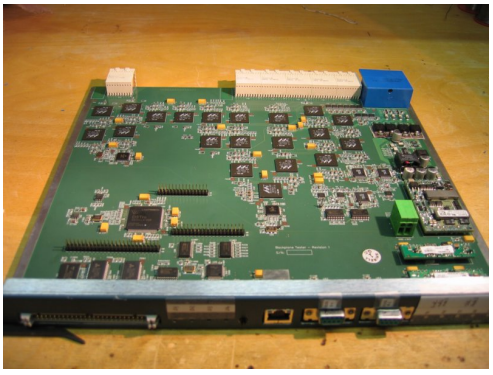

Fig. 4. View of Hub1 PCB

The key design challenge in this PCB design was the routing strategy for the 156 differential pairs related to Fabric Interface Channels and the routing of the 52 differential pairs related to Base Interface Channels.

III. BACKPLANE TESTER SOFTWARE ARCHITECTURE

The control of the AdvancedTCA backplane tester is achieved from an externally remote PC over two serial interfaces: a RS-232 serial link and an Ethernet interface. The serial port is used to load the software to the controller and through the Ethernet interface passes the communication between the applications operating on the controller and on a remote PC. The control of the backplane tester requires the distribution of the software between the remote PC and the Lantronix.

The software architecture was therefore based on the client-server model. The client application is running on the remote PC, to which the user interacts directly by defining the test to be performed on the AdvancedTCA backplane via a user interface created in an Excel spreadsheet. The tasks of the client are to take the user specifications, to transform and to send them to the server for processing. This is achieved by using scripting in Visual Basic and C programming language for the client and the server application.

The server application, previously downloaded via the RS-232 serial port from the remote PC, is stored into the serial flash on the Hub1 board. The server emulates a MDIO master interface in order to control the Alaska devices to generate the traffic over the backplane, and monitor their state registers for obtaining statistics about each differential pair involved in the test over the backplane. The client final task is to periodically fetch the monitoring statistics back from the server, to report them in real-time to the user and also to keep them in logs for later use.

The statistics contain information about the status and the

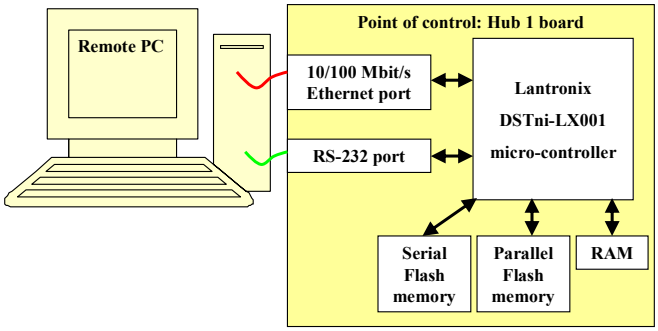

Fig. 5. Micro-controller Implementation

number of counted errors at the receiver of each differential pair tested over the AdvancedTCA backplane. The faulty differential pairs for a given test show up immediately in terms of connector row and column and slot position in the chassis.

IV. RESULTS

Using the boards of the AdvancedTCA backplane tester, we performed a series of BER tests by sending traffic for several hours under different conditions on all the 300 differential pairs in the dual-star fabric interface. The generated traffic was represented by various Jitter and PRBS patterns with different amplitudes and pre-emphasis levels for the transmitted signal. From a signal integrity point of view the most interesting was the Jitter mixed frequency test pattern (K28.5), extensively used in the further tests and measurements.

For the signal integrity evaluation and a cross-check of the software BER tests, we measured the received signal along the backplane. A Lecroy 5Ghz Serial Data Analyzer (SDA 5000A) [6] was used together with a 6 GHz differential probe (WaveLink Differential Probe D600A-AT) [7], equipment which would adequately cover up to the third harmonic in the 3.125 Gbps transmitted signal. The common method used for measuring the signal quality degradation across the backplane transmission channel is the eye-diagram resulting from the superposition of bit waveforms for all bits contained within the given PRBS pattern. Fig. 6 shows the eye-diagrams and their measured parameters at the receiver point corresponding to the longest distance in the AdvancedTCA chassis (Node board 14). The two diagrams are for transmitted signals of 700mV (Fig.6 a) and 1100mV (Fig 6 b)

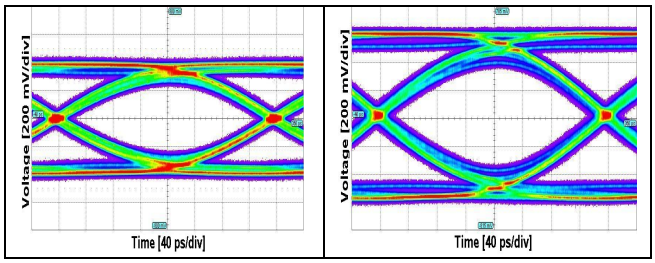

Fig. 6 AdvancedTCA Backplane Measurements (Pre-emphasis=0%)
a.) Transmitted Amplitude=700mV b.) Transmitted Amplitude=1100mV
Eye-opening=510 mV Eye-opening=740 mV
Eye-amplitude=690 mV Eye-amplitude=1 V

respectively.

The measured eye-opening exceeds even for the smallest transmitted signal amplitude, the value of 400mV mentioned in the XAUI specifications. These measurements confirm that the eye at the receiver has a more than sufficient opening, allowing for any Alaska device to reconstruct the received signal without any errors. Indeed, the BER tests performed for the same conditions detected no error on any of 300 differential pairs across the backplane.

The high-frequency losses in the backplane environment can be compensated for, and the eye-opening improved even further at the receiver by using a certain amount of pre-emphasis [8]. Further BER tests and measurements have been performed on the AdvancedTCA backplane for increasing levels of pre-emphasis and the improvement measured on the received signal. The effect of pre-emphasis added to the transmitted signal is shown in Fig.7, where the output from Marvell Alaska device with 0% pre-emphasis is shown in red and with 60% in black. The pre-emphasis effect is obtained by increasing (pre-emphasising) the amplitude of the single bit that follows the transition from one-to-zero or zero-to-one, while keeping untouched the amplitude of successive 'ones' and 'zeros' bits in the transmitted pattern. The nominal amplitude of both signals is 700mV with the exception of each bit period after a transition which is being driven to 1100mV, corresponding to 60% of pre-emphasis.

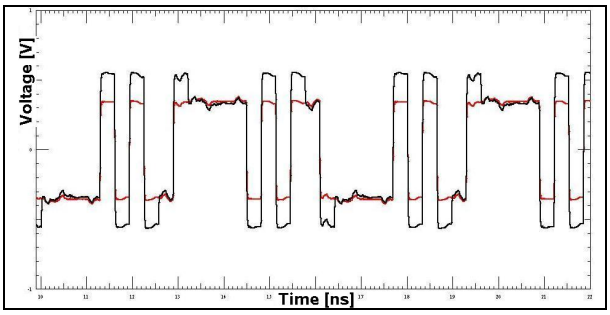

Fig. 7. Pre-emphasis Effect on the Transmitted Signal

Fig. 8 shows the eye-diagram measurements at the receiver of Node14 for the transmitted signals above.

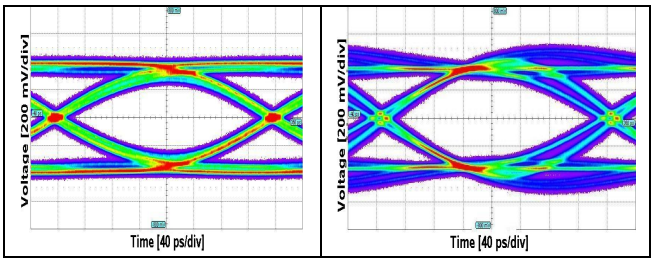

Fig. 8. AdvancedTCA Backplane Measurements

a.) Transmitted amplitude=700mV
Pre-emphasis=0%
Eye-opening=510 mV
Eye-amplitude=690 mV

b.)Transmitted amplitude=1100mV
Pre-emphasis=60%
Eye-opening=600 mV
Eye-amplitude=720 mV

We see here that for the same nominal bit amplitude we have improved the eye opening by 10% from 510mV to 600mV. However there remains overshoot from the pre-emphasised signal that indicates that too much pre-emphasis has been introduced leaving excessive voltage swing.

Different combinations of nominal voltage level and pre-emphasis were tried and the best results were obtained using 900mV signals with 33% pre-emphasis as shown in Fig 9. In this case an improvement of almost 20% is achieved with an eye opening of 740mV and no overshoots resulting in a clean and balanced eye-diagram. Thus although there were no errors observed for a signal without any pre-emphasis, we have shown that considerable improvements can be made using pre-emphasis to overcome any problems with production variations of backplane quality.

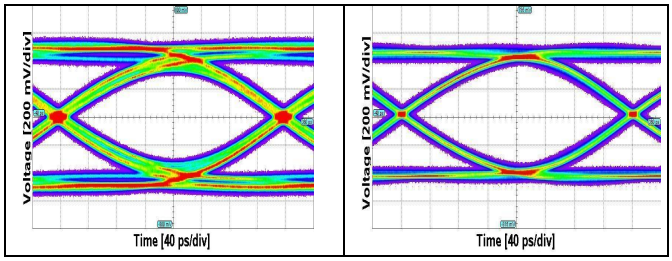

Fig. 9. AdvancedTCA Backplane Measurements

a.) Transmitted amplitude=900mV
Pre-emphasis=0%
Eye-opening=575 mV
Eye-amplitude=850 mV

b.)Transmitted amplitude=1100mV
Pre-emphasis=33%
Eye-opening=740 mV
Eye-amplitude=820 mV

The AdvancedTCA tester was also intended to be used to explore backplane behaviour under sub-optimal conditions. The Alaska device is actually programmed by defining the level of 'de-emphasis' which means that the swing of the post transition bit is set to a known value, and the nominal level of subsequent bits defined according to the chosen level of pre-emphasis. In the limit we can require up to 300% of pre-emphasis which for a first bit swing of 1100mV yields a nominal bit swing of only 366mV. This has the effect of reducing the eye-opening to much less than the nominal 400mV required by the XAUI standard. This is shown clearly in Fig 10 where the eye opening at the receiver has been thus artificially reduced to only 220mV together with an obvious signal overshoot.

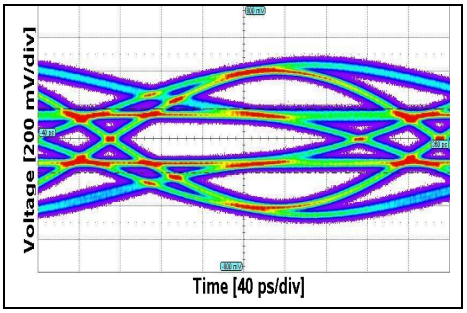

Fig. 10 AdvancedTCA Backplane Measurement
Transmitted amplitude=1100mV
Pre-emphasis=300%
Eye-opening=220 mV
Eye-opening=300 mV

Only under these extreme conditions we did start to observe any transmission errors on this particular AdvancedTCA backplane. Error free operation has been observed for eye openings as little as 450mV, which means that for this particular driver and receiver technology there is nearly a factor two of performance headroom since we can easily achieve an eye opening of 740 mV.

V. CONCLUSIONS

The AdvancedTCA Backplane Tester has been developed, built and exploited. We demonstrated its functionality by performing BER tests and measurements. The backplane under test was very well engineered and could sustain a good signal quality even for the minimum transmitted pre-emphasis over the longest distance in the chassis.

To obtain the quality estimation for any particular AdvancedTCA backplane we used the pre-emphasis circuitry of the Marvell Alaska technology to introduce artificially low thresholds of operation so that any significant PCB manufacturing anomalies would result in detectable error conditions. In practice the AdvancedTCA backplane performed within the requirements of the standard.

ACKNOWLEDGMENT

The authors would like to thank to Alexis Lestra for his major contribution in bringing the design through the production cycle to a successful prototype implementation, as well as his useful discussions and technical assistance during the development and testing of the system.

REFERENCES

[1] http://www.picmg.org/specdirectory.stm#_PICMG_3.0
[2] http://www.ist-esta.org/index.cfm?PID=56
[3] "88X2040 Datasheet – Integrated Single Chip Quad 3.125/3.1875 Gbps Transceiver", Rev. E, September 15, 2003
[4] "IEEE Draft P802.3ae /D5.0" – May 1, 2002, Annex 48A
[5] http://www.lantronix.com/pdf/DSTni-LX_PB.pdf
[6] http://www.lecroy.com/tm/products/Analyzers/home.asp?mseries=10
[7] http://www.lecroy.com/tm/products/Probes/Differential/WaveLink/default.asp
[8] "White Paper on Transmit Pre-Emphasis and Receive Equalization" – Johnny Zhang, Zhi Wong, , Mindspeed Technologies, October 31, 2002