An Integrated into FPGA System for Optical Link Testing and Parameters Tuning

Anton Kuzmin*, Dietmar Fey*, and Ulrich Lohmann†

*Friedrich-Alexander-Universität Erlangen-Nürnberg (FAU), Computer Architecture

{anton.kuzmin,dietmar.fey}@informatik.uni-erlangen.de

†Chair of Micro- and Nanophotonics, Fern University in Hagen

ulrich.lohmann@fernuni-hagen.de

Abstract—Development, characterization and performance optimization of systems utilizing FPGAs with high-speed serial transceivers to implement optical links with 1 to 10 Gbps data rate is a complex task and it poses several challenges for design engineers. In this paper, an effective approach is presented designed to address these challenges based on the use of diagnostic features implemented in the transceivers and a soft-IP microcontroller system instantiated in the FPGA. The use of the soft-IP controller allows a single-point access to the control and diagnostic interfaces of all components forming the link. Combined with computational capabilities and a high-level programming language interpreter running on the soft-IP CPU inside the FPGA, it enables extensive optical link performance evaluation without relying on any additional test and measurement equipment and significantly shortens debugging and testing times. Two generations of the system including hardware, soft-IP microcontroller system and embedded software are presented. The implementation demonstrates the feasibility and effectiveness of the proposed approach to utilization of on-chip diagnostic capabilities.

Index Terms—Optical fiber communication; Transceivers; FPGA; Microcontrollers; Embedded software

I. INTRODUCTION

Modern applications including rich media content transport, on-the-fly image processing, high bandwidth data acquisition for experimental physics, and high performance computing, require ever increasing serial communication data rates. At the same time, latency requirements remain strict and significantly limit possibilities for error correction and therefore call for a lower number of acceptable errors in the communication channel. FPGA devices with integrated high-speed serial transceivers and optical interconnects provide a very efficient and flexible platform for implementing such demanding applications and can be found in an increasing number of systems. A prototype system combining FPGA with optical interconnect implemented by the authors was presented at the VALID 2012 [1]. This paper demonstrates an extension of the original test system to a platform for reconfigurable computing with parallel optical links. Various examples and applications of optical interconnects could be found in [2]–[6].

One of the major challenges is a parameter tuning of the various components forming an interconnect to achieve the lowest possible probability of bit errors. The problem is that accurate measurements at low error probabilities require very long times even at high data rates to accumulate statistics

for a given confidence level while the parameter optimization space is relatively big. Additional complications arise from the fact that various components of the link have very different interfaces for setting parameters. In most cases, they are supported by proprietary tools with limited functionality for automatically tuning link parameters. The application of these tools often requires a connection of the system to external test and measurement equipment. The limitations associated with its usage become increasingly severe with a tighter integration between the FPGA and the optical transceiver blocks as recently proposed by Li et al. [7]. This level of integration makes electrical signals between the FPGA and optical transceiver practically inaccessible for external test equipment.

This paper presents an approach designed to address challenges associated with the testing, parameter tuning and performance monitoring of optical interconnects in FPGA-based systems. The approach is based on the use of a soft-IP controller embedded into the FPGA to perform two major tasks: link performance measurements and control of parameters of the different components forming the link.

The paper has the following structure. In the first section, an example of utilization of FPGA built-in transceiver diagnostic capabilities is presented and the key differences in the approach chosen by the authors are outlined. In the subsequent section, an overall inter-FPGA transceiver-based serial link structure is shown followed by brief description of its components and their respective configurable and tunable parameters. Then, a Bit Error Ratio (BER) [8] is introduced as an integral characteristic of link performance. An optimized algorithm for obtaining an accurate BER scan plot (bath-tub curve) is described. It can be used for indirect eye diagram width measurement by introducing a phase shift into a signal sampling point inside the receiver. The eye diagram width may serve as an indicator of the link performance and is used as a target function for the link parameter optimization.

Implementation aspects of the FPGA-based optical link test system are discussed in the next parts of the paper along with the obtained link performance measurement results. Comparison of the measured BER levels obtained on the prototype system for different optical modules confirms the validity of the implemented approach.

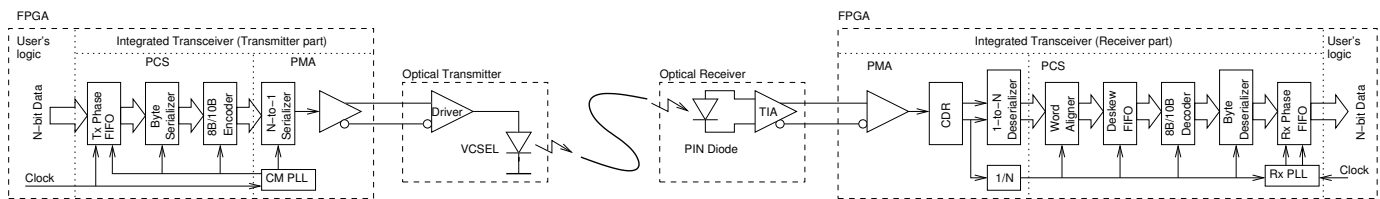Limitations of the developed prototype system are discussed

Fig. 1.   Simplified inter-FPGA serial optical link structure.

in the next section. The discussion is followed by a presentation of changes made to the prototype during development of the second generation of the system. The changes to all system components from hardware to test software are shown along with the reasons behind particular design decisions. A novel optical cable connector compatible with the flat optical cables used in the system is presented in the next section. The connector provides an efficient mean to implement a full mesh connection between twelve boards.

An application of the developed hardware and system components to construct a reconfigurable computing platform utilizing optical interconnects between the FPGAs is discussed in the concluding section.

II.   RELATED WORK

Usage of FPGA for testing communication channels has been previously described. For instance, in [6] an implementation of the Bit Error Ratio Tester (BERT) based on the Altera's Stratix II GX transceivers is presented and compared with a commercial stand alone tester. It is shown that the results obtained with the FPGA implementation comply with the results of the stand alone tester. However, the implementation still utilizes external equipment to control the test system and to collect the measurement results.

This paper, while similar in overall approach to the one proposed by Xiang et al. [6], presents notable improvements in several areas. The most important of these is the implementation of a functionally complete test system inside the FPGA. Additionally, the flexibility of the implemented system allows extension of its hardware and software components to support interfaces for monitoring and controlling the parameters of the various components of the link without external equipment. Another improvement presented in this paper is an adaptation of eye-width as a link performance indicator instead of a raw BER. The eye-width can be measured significantly faster at low bit error probabilities with the aid of diagnostic circuitries integrated into the transceivers and therefore is more efficient as a target function for the parameter space exploration and link performance optimization.

III.   OPTICAL LINK STRUCTURE

A block diagram of an optical digital communication link is shown in Figure 1. The link data path consists of a transmitter, an electro-optical converter (VCSEL with its driving circuits), an optical fiber, a photo detector (PIN diode and transimpedance amplifier) and a receiver. The transmitter and

receiver are further divided into a Physical Coding Sublayer (PCS) and a Physical Medium Attachment (PMA) sublayer.

The PCS blocks are responsible for byte serialization/deserialization, byte ordering, rate matching, and 8B/10B encoding/decoding. All these functions are essential for the implementation of a reliable digital data channel. However, in this work, we concentrate on the physical layer performance measurements leaving the problems related to the coding sublayer out of the scope of the research.

The transmitter part of the transceivers integrated into the FPGA allows the tuning and run-time changes of several parameters. Among them are clock multiplication phase-locked loop (PLL) dividers and bandwidth, output driver common mode voltage, differential voltage output swing and preemphasis aimed at reducing the negative effects of inter-symbol interference. The receiver part, in turn, has the following tunable blocks and parameters: on-chip termination, adaptive equalization, decision feedback equalization, receiver input common mode voltage and gain. These blocks have a crucial impact on the signal quality on the input of the Clock and Data Recovery (CDR) circuitry, but their influence cannot be measured directly because the signal after these stages is not physically available outside the chip and cannot be connected to external measurement equipment. The CDR block provides a built-in diagnostic support circuitry to facilitate assessment of the signal quality on its input.

The hardware interfaces, which are necessary to change all the transceiver's parameters and to access the diagnostic circuits, are available to the logic programmed into the FPGA. Chip and design software vendors provide tools to access these interfaces, however, their use requires a connection between the development workstation with CAD software and the FPGA. The electro-optical components of the link have their own sets of tunable and monitoring parameters, such as driver and receiver power levels, VCSEL modulation and offset currents, temperatures and thermal compensation coefficients, signal power detected at the receiver input, etc. Access to these features is implemented through another set of vendor-specific interfaces and also requires a development workstation with a connection to the target system. Such connections may be not feasible in the embedded system while access to the interfaces is still highly desirable or even required. This problem may be addressed by integration of IP cores for all required management interfaces into the system instantiated in the FPGA.

The flexibility of a soft-IP microcontroller system inside

the FPGA allows the implementation of a single-point access to the management interfaces of all the components forming the link. Combined with built-in link diagnostic capabilities controlled by the same microcontroller system it results in a complete test system that enables link performance testing and parameter tuning without relying on any external equipment. Additionally, it is available not only during development and testing of the system but also after its deployment.

IV. LINK PERFORMANCE INDICATORS

Two link operation quality indicators are introduced in this section along with a description of an algorithm used by the authors to measure "eye-width" with the transceiver's built-in diagnostic circuits.

A. Bit Error Ratio

The integral quality of operation of a serial link is characterized by its Bit Error Ratio (BER): a ratio of the number of bits received with errors to the total number of bits transmitted through the link: $BER = N_{err}/N$. This ratio is used for both measured and actual values. A BER is usually measured with a special piece of test equipment, so called Bit Error Ratio Tester. It consists of a data pattern generator, a reference quality receiver, a digital comparator and counters for transmitted bits and errors. The flexibility of an FPGA allows to implement all blocks of a bit error ratio tester in programmable logic in the FPGA itself.

If single bit errors in a serial link may be viewed as independent events and conditions do not change over time, then the actual BER value is a probability of a single bit error (p_e) and the measured value approaches the actual BER in the limit: $\lim_{N \to \infty} N_{err}/N = p_e$. It is not possible for BER measurement to transmit an infinite number of bits since it would require an infinite measurement time and a way to measure the BER with a given accuracy is required. For practical application it is often enough to know that the BER is below some threshold with a given confidence while its actual value is irrelevant. As the literature shows (for instance, in [9]), if more than N_0 bits were transferred during the test with no errors detected, then with probability α the actual BER is less than p_e:

$$N \geq N_0 = \frac{1}{p_e} \ln \frac{1}{1 - \alpha}$$

This number of bits (N_0) sets a lower limit on the test duration when no errors are observed. At a data rate of 5 Gbps it takes approximately 10 minutes to reach a 95% confidence that BER is lower than 10^{-12}, for the BER level of 10^{-15} it would require almost a week. The long runtime required makes it impractical to use the BER directly as a target function for the link parameters optimization. It would take enormous amount of time to find an optimum in the parameter space even if only a small fraction of all possible parameter combinations yielded a bit error ratio lower than 10^{-12}.

B. Eye-Width and its Measurement

The quality of a signal may be analyzed by evaluating its eye diagram: a picture on an oscilloscope display resulting from observing a transmission of a pseudo-random binary sequence with properties representative of the physical layer encoding used in the link. The width and height of an opening of the central part of the diagram ("eye") serve as indicators of the signal quality and may be used as target functions for the link parameter tuning. However, the signal on the input of the receiver CDR unit is not available for direct measurements. Therefore built-in diagnostic circuitries of the receiver should be utilized.

Serial transceivers integrated into the Altera Stratix IV GX FPGAs include special circuitry that facilitates measurements of the eye opening on the input of the CDR block [10]. The circuitry allows shifting of a sampling point of the signal from its optimal position in the center of the unit interval (UI) under external control. Then bit error ratio is measured for each phase offset. For sampling points close to the center of the eye opening, there will be no significant increase in the bit error ratio. For sampling points closer to the signal slopes the number of observed errors will gradually increase. Finally, in the area of the signal edge crossing widened by a jitter, a receiver will not be able to achieve synchronization with its input signal resulting in the observed bit error ratio of 0.5. From these measurements of the BER at signal sampling points distributed through the UI the eye opening and jitter characteristics of the signal may be deduced [9].

The key benefit of this approach is that the conclusion regarding the signal quality and, therefore, link parameters, may be reached by a number of BER measurements with different phase offsets through the UI instead of one at the optimal sampling point. However, each of these measurements needs to achieve a given confidence level at a much higher target BER and, therefore, requires significantly shorter runtime.

An algorithm implementing this approach can be further optimized to reduce the number of required BER measurements at the center of the eye opening, where the bit error ratio is low. These measurements take up most of the time and effectively provide no useful information. Several approaches to such optimization are described in [9].

Figure 2 illustrates the behavior of the modified algorithm implemented by the authors and shows an eye-diagram reconstructed from the measurements. As a first step (marked with 1 in the figure) an initial scan through the entire unit interval is performed with high target BER (10^{-7}). From these measurements, an approximate location of the eye boundaries is determined. At the second stage the BER is measured at the center of the eye opening to make sure that the target BER level (10^{-12}) is achievable at the close-to-optimal sampling point (2). Then, the BER is measured at sample points from the eye opening boundaries detected during the first scan towards the center to determine points where the target BER level is achieved (3). The distance between these points (eye-width) serves as a measure of the signal quality at the input of the
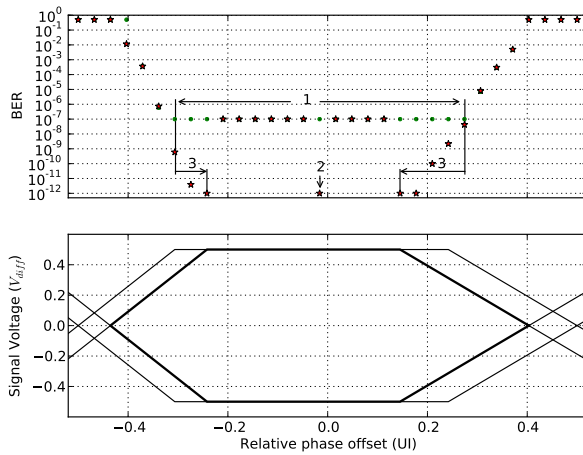
Fig. 2.   "Bath-tub" curve scan algorithm and reconstructed eye diagram.

Fig. 3.   Experimental system and loopback configurations.

receiver CDR unit and may be used as a target function for the link parameters' tuning.

The described algorithm for eye-width measurements reduces the number of BER samplings within the eye opening. For the diagram shown in Figure 2, it took only 55 minutes to collect all the data. An exhaustive UI scan under the same conditions takes 150 minutes but provides no additional information on the link operation.

V. PROTOTYPE SYSTEM IMPLEMENTATION

To confirm the usefulness of the approach described to the optical link testing and parameter tuning and to create a base set of tools and building blocks to be used in future projects the authors implemented a prototype system. The system consists of hardware, a set of IP blocks, embedded software and development tools and facilitates debugging, testing and evaluation of the components. A photo of the assembled system hardware is shown in Figure 3 and components of the system are described in the following sections.

A. Hardware Platform

The system is based on the Altera Stratix IV GX FPGA (EP4SGX230KF40C2) installed on a TerasIC DE4 board. Through an adapter board with SMA connectors and a set of coaxial cables the DE4 board is connected to SFP+ evaluation boards hosting optical transceiver modules. Hot-pluggable SFP+ transceivers used in the system provide duplex LC-type optical connectors for the Multi-Mode Fiber. Management interface of the transceiver modules (I^2C) is accessible from the FPGA and is used for the monitoring of their parameters.

The highly modular construction of the hardware platform enables experimentation with different components and link configurations. During development and validation of the system several loopback configurations were used as shown in the diagram on Figure 3. The shortest possible one is an electrical loopback connecting the FPGA transmitter output signals directly to the input of the receiver (1). The second
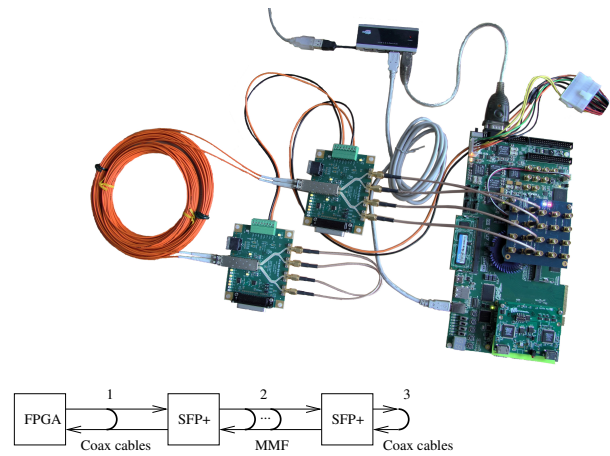
tested configuration uses a single optical transceiver with its input and output connected via a Multi-Mode Fiber (MMF) loopback (2). The length of the fiber loop used in the tests ranged from 15 cm to 15 meters. This loopback configuration is the closest to an actual optical link where the signal passes through one electro-optical and one opto-electrical conversion and a single fiber segment.

The most elaborate loopback configuration tested utilizes two transceiver modules and an electrical loopback on the "remote" side of a duplex fiber link (3). While this link exceeds configurations, which would be found in practical applications it is still interesting as it allows an easier separation of influence on the signal quality from different components of the link and serves as a model of a less favorable environment with longer links and a higher number of interconnects along the signal path.

The transceivers available in Stratix IV GX FPGA provide an on-die scope capable of $1/32$ unit interval resolution at data rates up to 6.5 Gbps [10]. Comparable technology is available in the transceivers integrated into the Xilinx Virtex-6 FPGA family. As an additional feature these transceivers are capable of a vertical scan of an eye-diagram [11], however, this functionality has not yet been explored by the authors so far.

B. System-on-Programmable Chip and IP Cores

The architecture of a soft-IP microcontroller system instantiated in the FPGA is shown in Figure 4. The system consists of the following main blocks: NIOS II CPU core with a small on-chip ROM containing boot code, a controller for external SRAM and FLASH, UART for communication with a control terminal, cores for the test pattern generator and checker, interfaces to access the transceiver configuration and diagnostic features, I^2C master cores for connection to the management interface of the SFP+ modules. The entire system utilizes only a small fraction of the available FPGA resources: the logic utilization is 3%, and available memory and DSP blocks are used for less than 1%.
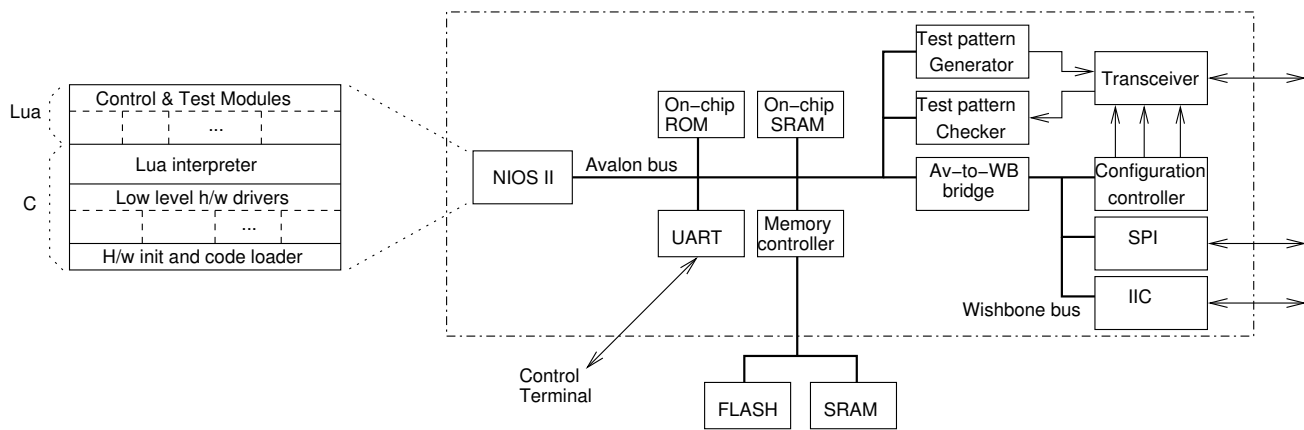
Fig. 4. Test System-on-Programmable Chip (SoPC) architecture.

The IP cores forming the system were taken from three sources. The first one is the library supplied by the FPGA vendor (Altera in this case). The cores are optimized for a specific FPGA architecture, but no source code is provided and the cores are not available on FPGAs from other vendors. The second source of IP cores for the system is a collection of free and open cores hosted on the OpenCores site [12]. These cores are provided under free licenses and their source code is available. This makes it possible to implement these cores in systems on different FPGA architectures. The price for such flexibility is the time and effort required for integration and adaptation, and the required time and effort is generally greater than for FPGA vendor supplied IP cores.

These two sources of IP cores, while covering most of the functionality, still do not provide several crucial interfaces required in order to access transceiver configuration and diagnostic interfaces. These missing parts were created by the authors by means of custom HDL development as the third source of IP blocks, and this required most effort.

Since the IP cores from different sources have different interfaces their integration into a working system is a technical problem in itself and required the development of "adapter" modules. The two primary on-chip interconnects used in the system are Avalon [13] and WISHBONE [14].

Overall, a combination of the readily available blocks (both proprietary and free) and those developed in-house proved to provide a reasonable and time efficient way of implementing the prototype system.

C. Embedded Software

The monitoring and control of all blocks forming the optical link, BER testing and processing of the test results are handled by an embedded software running on the NIOS II soft-IP CPU instantiated in the FPGA.

Low level software to access all hardware interfaces is implemented in the C programming language and its functionality is made available to the Lua interpreter. Lua, as is stated on its web-site [15], "is a powerful, fast, lightweight, embeddable scripting language". These properties make it very

attractive for a wide range of applications including game development, mobile devices and embedded software [16]. A tight integration with C and an interactive interpreter facilitate an efficient development of diagnostic, testing and debugging software for embedded hardware systems.

Availability of an ANSI C compiler and a basic C run-time library are the only requirements to port Lua to a new platform and it was extremely easy to get an early prototype running on NIOS II. The efforts invested in the porting and support of Lua interpreter on the soft-IP microcontroller system in the FPGA were rewarded in the flexibility of the resulting system and increased development productivity.

Access to the interactive environment is very useful during embedded hardware development and debugging as it saves a lot of time in the edit-compile-load-run development cycle. Since the "hardware" itself is a soft-IP system instantiated in the FPGA this time saving becomes even more important: on the one hand, the system is malleable and experimental and includes design errors, on the other hand, traditional software development cycle is complicated by a separate FPGA design flow with longer iterations. With this additional complexity an availability of tools facilitating quick experiments and tests running directly on the target platform is a key factor for effective development. Our experience shows that Lua fits this role perfectly and allows rapid localization of the design errors both on the hardware and software levels. All the link configuration and BER measurement software in the system are implemented as a set of Lua modules.

VI. MEASUREMENT RESULTS

Measurements on the test system were performed for data rates in a range from 1 to 5 Gbps with various loopback configurations. The SFP+ module used in most experiments is the Avago AFBR-703SDDZ. The module is capable of data rates up to 10 Gbps and, as expected, performs excellently in the tested data rate range. Even with the most demanding loopback configuration the eye diagram opening for the 10^{-12} BER level is approximately 40% (80 ps) of the unit interval (200 ps at 5 Gbps).
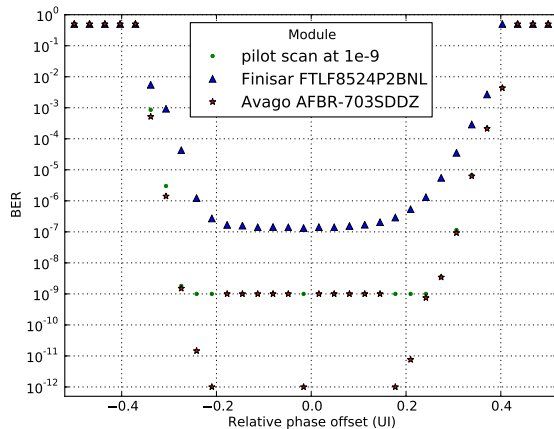
Fig. 5. Comparison of "bath-tub" curves for two SFP modules at 5 Gbps.

Several data patterns with different spectral characteristics were used in the experiments. Two test patterns that specifically check the link performance at the edges of its frequency band are the Low Frequency (LF) and High Frequency (HF) patterns. The other test patterns are Pseudo-Random Binary Sequences (PRBSx) generated by a linear feedback shift register with the length x. The lengths of 7, 15, 23, and 31 bit were used. The test results show slight dependency on the data pattern used, however, detailed analyses of this dependency have not yet been yet performed.

To validate the test system and confirm that the measurement results adequately represent link quality an SFP module with a lower maximum data rate has been used: Finisar FTLF8524P2BNL. According to its documentation the module is capable of data rates up to 4.25 Gbps. Experiments show that up to this limit it demonstrates BER $\leq 10^{-12}$, also the eye width is smaller than that with the Avago module. The bathtub scan results for both modules at 5 Gbps are shown in Figure 5. This data rate is outside of the specified range for the Finisar module and this is clearly visible from the diagram: even in the vicinity of the ideal sampling point BER does not achieve 10^{-7} level.

The results obtained allow the conclusion that the developed test system provides reliable data on the optical link performance and may be used to compare different link implementations and to tune parameters of the link. The comparison of the measurement results obtained with different data patterns may provide additional information that could be useful for optimizing link performance.

VII. SECOND GENERATION OF THE SYSTEM

While the developed system proved the feasibility of the selected approach and served as a convenient platform for experiments with the optical links and IP cores, it also had severe limitations, which made further usage of the system and its components problematic. The following sections describe these limitations along with the changes implemented by the authors in the second generation of the system to overcome the discovered limitations.

A. Hardware Platform

The most obvious problems of the prototype platform at the hardware level were a limited number of the supported optical channels and a low integration level leading to a number of separate boards interconnected with a web of cables. While being beneficial at the early stage of the project enabling fast system setup and reconfiguration times, this approach does not scale well beyond simple desktop setup with one optical link. The most elaborate system configuration theoretically achievable with this approach would have eight optical links and require eleven boards for loopback configuration only. The situation would be even worse considering the interconnection of several FPGA boards. The other limitation of the prototype hardware platform is the unavailability of the optical transceivers' current consumption monitoring.

Recent advances in the optical transceiver technology enable higher level of integration and power efficiency than achievable with the SFP+ modules used in the prototype system [17]. To benefit from these improvements in the technology the authors designed an add-on card for the DE4 FPGA board. The add-on card mounted on top of the DE4 FPGA board with flat optical cables connected is shown in Figure 6. The card hosts Avago transmitter and receiver MiniPOD™ modules and connects them to the FPGA serial transceivers making use of all twelve channels available on the DE4 board extension connectors. The transmitter and receiver configuration and monitoring interfaces are connected to the FPGA as well. In addition to the management functionality integrated into the modules, the board implements circuitry to individually monitor power consumption of the modules on all power supply rails.

The developed add-on card makes twelve duplex optical links available for the FPGA without any additional boards or electrical cables connected. The transceiver and receiver modules make use of a flat ribbon optical cable to implement a high-density optical connection.

Two FPGA boards may be directly linked with ribbon optical cables to implement twelve duplex communication links. Extending the system to higher number of nodes would require more complex optical interconnection. One specifically important type of the interconnect is a full-mesh network where each node has a direct link to every other. A novel matrix optical connector presented in [18] implements a crossbar interconnect for twelve nodes.

The realization of an 4×4 crossbar interconnect is illustrated in Figure 7. A connector uses two fiber-matrix plates that are rotated by $90°$. Each 2D fiber-matrix plate combines four 1D fiber bundles. Because of the $90°$ rotation of the matrix plates, columns at the input side are connected to rows at the output side. Each output fiber bundle is connected to every input layer. Therefore, all the combinations of inputs and outputs are realized as required by the crossbar scheme.
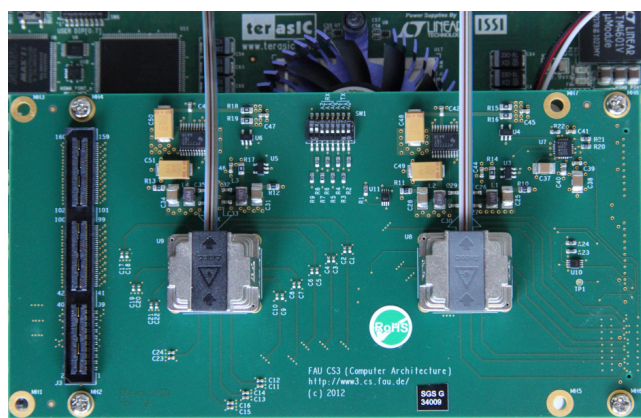
Fig. 6.   MiniPOD extension board mounted on top of DE4.

Despite the simplicity of the interconnection scheme, the realization of the fiber-matrix poses significant challenges. In Figure 8 a 12×12 connector is shown. It includes a matrix with 12×12 multimode fibers, each with a core diameter of $50~\mu$m. The distance between the single fibers is $250~\mu$m and the total area of the matrix is 2.85×2.85 mm^2. The difficulty in making large connectors is given by the requirement to drill high number of holes with tight position tolerances. A detailed analysis of the fabrication procedure may be found in [19]. This crossbar optical connector with 144 channels in total is available as a so called CrossCon®-device from the company EUROMICRON.
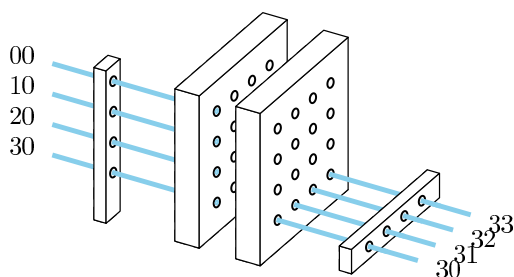

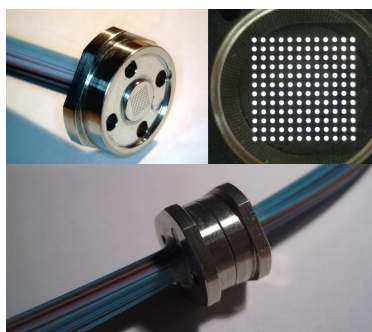Fig. 7.   Schematic of the novel 3D fiber optical crossbar approach

Fig. 8.   a) Prototype of a 12×12 fiber connector. b) Picture of the fiber matrix with 144 channels. The position tolerance of the holes is $\pm 3~\mu$m. c) Resulting crossbar interconnection.

B. System-on-Programmable Chip and IP Cores

Several factors limit a wider application of the System-on-Programmable-Chip solutions. One of the most critical is the utilization of FPGA vendor specific IP cores. To use the test system on an FPGA from a different vendor these blocks should be replaced with their functional equivalents available on the other platform, but supporting different system variants would increase the effort required. A more efficient approach is to replace the vendor specific IP cores with free and open-source equivalents available on different target platforms.

The most complex and important block in the prototype system specific to the Altera platform is the NIOS II CPU core, so the authors decided to replace it with one of the free and open-source CPU cores available. Several cores have been considered as a replacement candidates and the choice was made for LEON3 processor from Aeroflex Gaisler AB [20]. The LEON processor is "a 32-bit synthesisable processor core based on the SPARC V8 architecture. The core is highly configurable, and particularly suitable for system-on-a-chip (SOC) designs." The LEON3 core is distributed as a part of GRLIB IP library. The library also includes many communication peripheral controllers, configuration utility, and several pre-configured design examples. The configuration utility allows setup of various LEON3 core parameters (such as size of a register window, cache type and sizes, etc.) and configure system peripheral devices. At last but not least the system include Debug Support Unit (DSU), which allows external connection to the system via various interfaces (JTAG, serial, ethernet network) and on-chip software debugging access. Availability of this functionality is invaluable during the early firmware porting.

The remaining proprietary cores (test data pattern generator and checker, external bus controller) are easier to replace and do not require toolchain and embedded software porting efforts. The work on their replacement with equivalents available as open-source or developed by the authors is underway.

The modified architecture of the System-on-Programmable Chip is presented in Figure 9. Other than processor core replacement and associated changes in the peripheral devices and connection bus architecture there is one other notable change compared to the prototype system. To reduce FPGA resource usage and power consumption associated with high-speed circuitries and to simplify time-closure of this part of the design without compromising the ability to test all twelve serial communication links the decision has been made to provide just one data pattern generator and checker and implement a daisy-chain connection for the serial transceivers on their digital data interface to the FPGA. In this configuration test data received on the first transmitter channel (which are the data sent on this channel and returned via optical loopback) are retransmitted on the second channel and so forth. A parallel multiplexer allows to select a transceiver channel to be connected to the data pattern checker.

The synthesis results for the system with the new processor core and modified architecture show that the FPGA resource
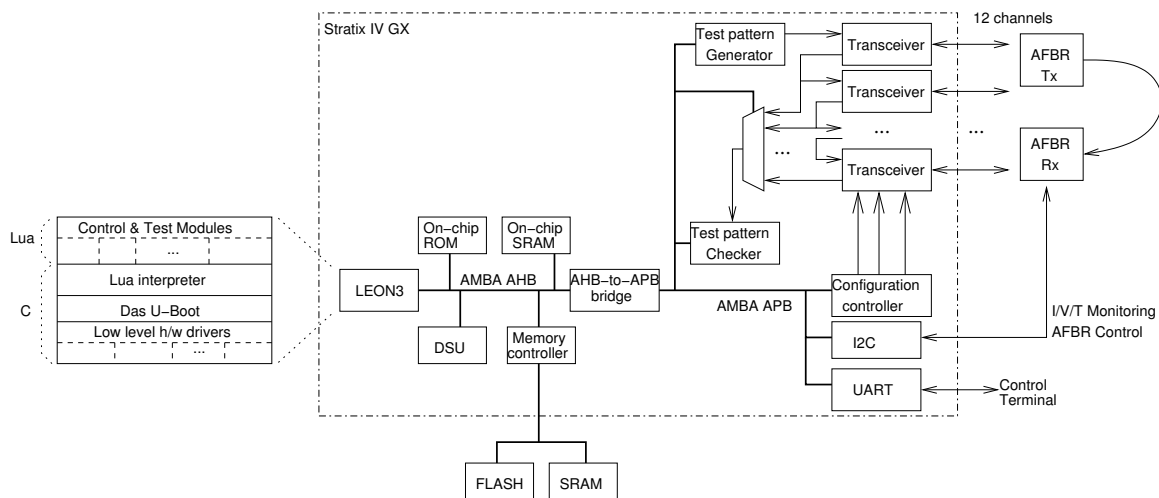
Fig. 9.   Second generation of the SoPC architecture.

usage has not been significantly affected by the changes and remains below 5%. That makes it possible to keep the test system in the FPGA configurations for different applications and to use the embedded LEON3 processor for management, debugging support and performance monitoring tasks in the application system.

C. Embedded Software

The chosen approach to implement embedded software for the SoPC on the interpreted language facilitating fast and interactive development directly on the target system proved to be an efficient one. That is especially notable in the areas where a non-standard interfaces or functionality have to be implemented or there is a need to experiment with different algorithms or test various approaches with a quick prototype code.

On the other hand, this approach requires to re-implement a lot of common support functions (e.g., filesystem and networking), which come for granted with a "standard" embedded software development tools and frameworks. The switch to a advanced processor core also requires additional efforts to be invested in development of the general purpose utility code such as processor boot, cache and interrupt management, etc. Anticipating utilization of the LEON3 processor inside the system for other tasks and the need for software support for a number of standard interfaces and protocols when the boards are integrated into a system the authors decided to introduce an extra level into the system embedded software.

This new level of software is to replace custom written processor bootstrap code, initialization and drivers for generic hardware interfaces and basic C runtime library from the prototype system. An industry standard open source firmware for the embedded platforms with a wide range of supported processor architectures currently is "Das U-Boot" [21]. It implements all low-level processor and interface management tasks and allow load and debugging of application software from different medias.

While the primary goal of U-Boot in an embedded system is to load application software (often it is a Linux kernel along with the initial ramdisk image) from external storage or over a network interface and pass control to it, the firmware provides and Application Programming Interface (API) for so-called "U-Boot standalone applications". These applications are loaded dynamically and can have an access to the U-Boot console I/O functions, memory allocation and interrupt services. The Lua interpreter with serial links and system control, monitoring and test software plays the role of such standalone application.

The modified embedded software architecture allows to combine positive sides of the interactive interpreter with custom experimental software running on the target system with an industrial-strength support for multiple interfaces, protocols and debugging capabilities coming with a standard cross-compiled firmware.

VIII. CONCLUSION AND FUTURE WORK

The implementation of the optical link test system clearly demonstrated the feasibility and effectiveness of the proposed approach to utilization of the on-chip diagnostic capabilities of FPGAs with high-speed serial transceivers. The use of the soft-IP controller instantiated in the FPGA allows a single-point access to the control and diagnostic interfaces of all components forming the link. Combined with computational capabilities and a high-level programming language interpreter running inside the FPGA, it enables extensive optical link performance evaluation without relying on any additional test and measurement equipment and significantly shortens the system debugging and testing times. As an additional benefit all the implemented functionality is still available in the deployed system and may be used for remote monitoring and diagnostics.

Detailed analysis of the dependencies between the test loopback configurations, data patterns, transceiver parameters and observed eye-diagram is required to develop effective
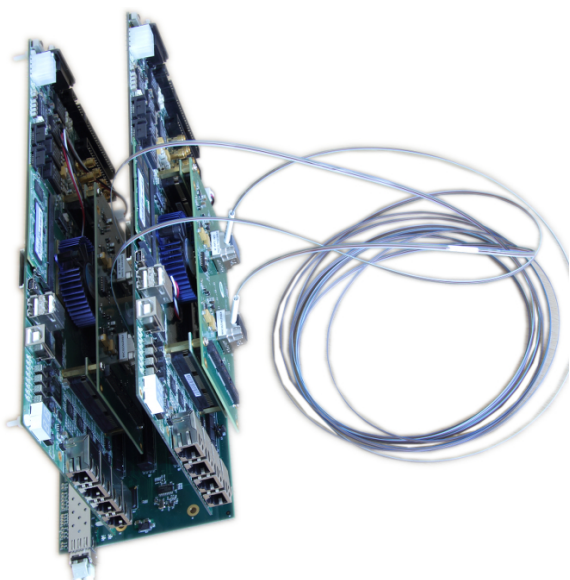
Fig. 10. Two DE4 boards cross-connected with parallel optical links.

algorithms for link parameters tuning. Current consumption monitoring hardware integrated into the system facilitates measurements of the system power efficiency. This work provides efficient tools for these researches and demonstrates their feasibility.

Another area for improvement is the automated integration of separate IP blocks from different sources into a system. Vendor specific tools have progressed notably in this area in recent years, however, they are still limited with regard to support of "foreign" IP cores. On the other hand, while efforts have being made to provide similar functionality for free and open-source cores, the tools that have emerged so far are not well integrated in the FPGA and embedded software design flows. The authors' experience gained in course of conversion the system to open IP cores shows that despite availability of automated integration and configuration tools a manual intervention and hand-written RTL code are still required to combine IP blocks from different sources into a working system.

The developed hardware platform, IP blocks and embedded software form a base for integration of parallel optical links into a multi-FPGA reconfigurable computing system. Current hardware platform consisting of two DE4 FPGA boards cross-connected by twelve parallel optical links is shown in Figure 10. This platform is used for development of streaming video processing and HPC applications. When a framework for such application reaches some maturity and its hardware resource requirements exceed limits of the current platform, the next extension step is to switch from dual-board system to the multiple cross-connected boards configuration.

REFERENCES

[1] A. Kuzmin and D. Fey, "Optical link testing and parameters tuning with a test system fully integrated into FPGA," in *The Fourth International Conference on Advances in System Testing and Validation Lifecycle (VALID 2012).* IARIA, 2012, pp. 121–126.

[2] A. F. Benner, M. Ignatowski, J. A. Kash, D. M. Kuchta, and M. B. Ritter, "Exploitation of optical interconnects in future server architectures," *IBM Journal of Research & Development*, vol. 49, no. 4/5, p. 755, July/September 2005.

[3] S. Nakagawa, Y. Taira, H. Numata, K. Kobayashi, K. Terada, and M. Fukui, "High-bandwidth, chip-based optical interconnects on waveguide-integrated SLC for optical off-chip I/O," in *Electronic Components and Technology Conference*, 2009, pp. 2086–2091.

[4] B. E. Lemoff, M. E. Ali, G. Panotopoulos, E. de Groot, G. M. Flower, G. H. Rankin, A. J. Schmit, K. D. Djordjev, M. R. T. Tan, W. Gong, R. P. Tella, B. Law, and D. W. Dolfi, "Parallel-WDM for multi-Tb/s optical interconnects," in *Lasers and Electro-Optics Society (LEOS) IEEE Meeting.* Agilent Technologies Laboratories, 2005, pp. 359–360.

[5] O. Liboiron-Ladouceur, H. Wang, A. S. Garg, and K. Bergman, "Low-power, transparent optical network interface for high bandwidth off-chip interconnects," *Optics Express*, vol. 17, pp. 6550–6561, 2009.

[6] A. C. Xiang, T. Cao, D. Gong, S. Hou, C. Liu, T. Liu, D.-S. Su, P.-K. Teng, and J. Ye, "High-speed serial optical link test bench using FPGA with embedded transceivers," in *Topical Workshop on Electronics for Particle Physics (TWEPP)*, 2009, pp. 471–475.

[7] M. P. Li, J. Martinez, and D. Vaughan. Transferring high-speed data over long distances with combined FPGA and multichannel optical modules. [Online]. Available: http://www.altera.com/literature/wp/wp-01177-AV02-3383EN-optical-module.pdf [retrieved: May, 2014].

[8] G. Breed, "Bit error rate: fundamental concepts and measurement issues," *High Frequency Electronics*, pp. 46–48, January 2003.

[9] M. Müller, R. Stephens, and R. McHugh, "Total jitter measurement at low probability levels, using optimized BERT scan method," in *DesignCon.* Agilent Technologies, 2005.

[10] W. Ding, M. Pan, T. Tran, W. Wong, S. Shumarayev, M. Peng Li, and D. Chow, "An on-die scope based on a 40-nm process FPGA transceiver," in *DesignCon.* Altera Corporation, 2010.

[11] Virtex-6 FPGA GTX transceivers user guide. UG366 (v2.6). Xilinx, Inc. [Online]. Available: http://www.xilinx.com/support/documentation/user_guides/ug366.pdf [retrieved: May, 2014].

[12] [Online]. Available: http://opencores.org/ [retrieved: May, 2014].

[13] Avalon interface specifications. [Online]. Available: http://www.altera.com/literature/manual/mnl_avalon_spec.pdf [retrieved: May, 2014].

[14] Wishbone B4. WISHBONE System-on-Chip (SoC) interconnection architecture for portable IP cores. [Online]. Available: http://cdn.opencores.org/downloads/wbspec_b4.pdf [retrieved: May, 2014].

[15] [Online]. Available: http://www.lua.org/about.html [retrieved: May, 2014].

[16] R. Ierusalimschy, *Programming in Lua.* Lua.org, 2006.

[17] Micropod and minipod 120g/150g/168g transmitters/receivers. [Online]. Available: http://www.avagotech.com/pages/en/fiber_optics/parallel_optics/minipod_micropod [retrieved: May, 2014].

[18] U. Lohmann, J. Jahns, A. Kuzmin, and D. Fey, "Optical multi-Gbps board-to-board interconnection with integrated FPGA-based diagnostics," in *Optical Interconnects Conference.* IEEE, 2013, pp. 120–121.

[19] U. Lohmann, J. Jahns, T. Wagner, and C. Werner, "Ultra-precision fabrication of high density microoptical backbone interconnection for data center and mobile application," *Proc. SPIE Optics and Photonics*, 2012.

[20] [Online]. Available: http://www.gaisler.com/index.php/products/processors/leon3 [retrieved: May, 2014].

[21] Das U-Boot – the universal boot loader. [Online]. Available: http://www.denx.de/wiki/U-Boot/WebHome [retrieved: May, 2014].