



European Coordination for Accelerator Research and Development

PUBLICATION

Interfaces and Communication Protocols in ATCA-Based LLRF Control Systems

Makowski, D (Technical University of Lodz) *et al*

04 February 2010

IEEE TRANSACTIONS ON NUCLEAR SCIENCE

The research leading to these results has received funding from the European Commission under the FP7 Research Infrastructures project EuCARD, grant agreement no. 227579.

This work is part of EuCARD Work Package **10: SC RF technology for higher intensity proton accelerators and higher energy electron linacs.**

The electronic version of this EuCARD Publication is available via the EuCARD web site <<http://cern.ch/eucard>> or on the CERN Document Server at the following URL :
<<http://cdsweb.cern.ch/record/1237832>>

Interfaces and Communication Protocols in ATCA-Based LLRF Control Systems

Dariusz Makowski, *Member, IEEE*, Waldemar Koprek, Tomasz Jeżyński, Adam Piotrowski, *Member, IEEE*, Grzegorz Jabłoński, Wojciech Jałmużna, and Stefan Simrock

Abstract—Linear accelerators driving Free Electron Lasers (FELs), such as the Free Electron Laser in Hamburg (FLASH) or the X-ray Free Electron Laser (XFEL), require sophisticated Low Level Radio Frequency (LLRF) control systems. The controller of the LLRF system should stabilize the phase and amplitude of the field in accelerating modules below 0.02% of the amplitude and 0.01 degree for phase tolerances to produce an ultra stable electron beam that meets the required conditions for Self-Amplified Spontaneous Emission (SASE). Since the LLRF system for the XFEL must be in operation for the next 20 years, it should be reliable, reproducible and upgradeable. Having in mind all requirements of the LLRF control system, the Advanced Telecommunications Computing Architecture (ATCA) has been chosen to build a prototype of the LLRF system for the FLASH accelerator that is able to supervise 32 cavities of one RF station. The LLRF controller takes advantage of features offered by the ATCA standard. The LLRF system consists of a few ATCA carrier blades, Rear Transition Modules (RTM) and several Advanced Mezzanine Cards (AMCs) that provide all necessary digital and analog hardware components. The distributed hardware of the LLRF system requires a number of communication links that should provide different latencies, bandwidths and protocols. The paper presents the general view of the ATCA-based LLRF system, discusses requirements and proposes an application for various interfaces and protocols in the distributed LLRF control system.

Index Terms—Accelerator control systems, accelerator instrumentation, accelerator RF systems.

I. INTRODUCTION

A POWERFUL digital Low Level Radio Frequency (LLRF) system is required to fulfil heavy demands of the control system for reliable operation of linear accelerators such as the FLASH or XFEL [1]. A real-time soft controller with a digital fast feedback and adaptive feed-forward is used to stabilize the cavity fields in accelerating modules [2]. A block diagram of the LLRF controller used for the FLASH accelerator is presented in Fig. 1. The LLRF controller submodule measures an electric field in cavities (transmitted, forward and reflected power) and

Manuscript received February 27, 2009; revised June 04, 2009. Current version published October 07, 2009. The research leading to these results has received funding from the European Commission under the EuCARD FP7 Research Infrastructures grant agreement no. 227579. The author is a scholarship holder of project entitled “Innovative education...” supported by European Social Fund.

D. Makowski, A. Piotrowski, G. Jabłoński, and W. Jałmużna are with the Department of Microelectronics and Computer Science, Technical University of Łódź, 90-924 Łódź, Poland (e-mail: dmakow@dmcs.p.lodz.pl).

W. Koprek, T. Jeżyński, and S. Simrock are with the Deutsche Elektronen-Synchrotron (DESY), D-22603 Hamburg, Germany.

Color versions of one or more of the figures in this paper are available online at <http://ieeexplore.ieee.org>.

Digital Object Identifier 10.1109/TNS.2009.2027234

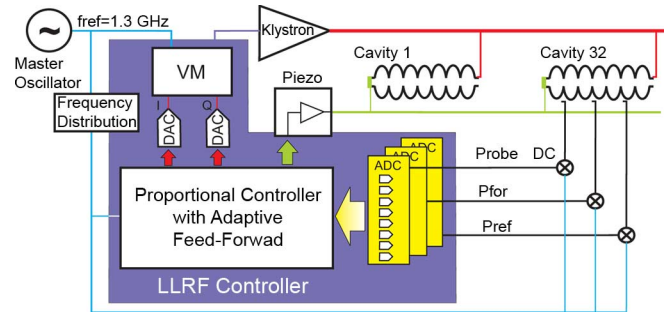


Fig. 1. A block diagram of the LLRF control system. The signal from the Master Oscillator with frequency $f = 1.3$ GHz is modulated by the Vector Modulator (VM), amplified by the klystron and forwarded to cavities. Cavity signals are converted from the cavity frequency to the intermediate frequency ($f = 54$ MHz) by Down Converters (DC), digitalized and used as input signals for the digital controller.

generates the complex control signal (Imaginary and Quadrature components) that is used to modulate the reference signal from the Master Oscillator. The control signal, amplified by a klystron, is distributed to the cavities through a wave-guide system. Accelerating cavities are supplied with 1.3 GHz modulated signal produced by 5 or 10 MW klystrons. Digital data processing is performed in the control system to produce the driving signal for the klystron.

The LLRF system, based on SimCon 3.1L [3] that is currently installed in FLASH, controls one cryogenic module comprised of eight superconductive TESLA cavities [4]. The LLRF system of FLASH is based on the VME (Versa Module Eurocard) architecture [5], [6]. Since the XFEL accelerator needs almost 1000 cavities, the LLRF system will consist of more than 32 RF stations. Therefore, it is desired to design a LLRF system able to supervise 32 cavities, one RF station comprised of four cryo-modules spread out over 50 m [7]. The LLRF controller supervising 32 cavities is connected to the other accelerator components with a significant number of analog and digital signals, e.g., 96 analog cavity signals, 32 analog and 32 digital signals for the fast and slow piezo tuners, 10 digital interlock signals, main RF input and output signals, and reference frequency and trigger signals [8]. In addition to operational demands, the system should be characterized with a high reliability, availability and modular design. ATCA and AMC modular standards offer a significant number of features helpful to a highly reliable LLRF system design [9], [10]. These standards offer hot-swap functionality and support redundancy for the most critical subsystems, namely: power supply, management, diagnostics and communication links. On the other hand, the number of required

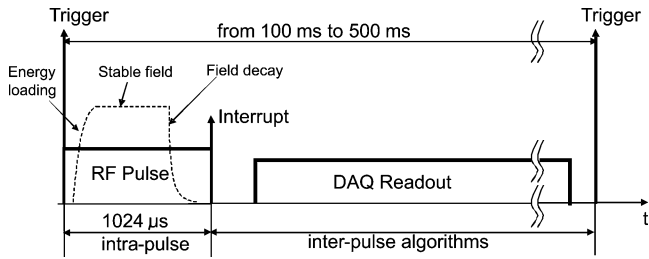


Fig. 2. The relationship of real-time algorithms with a single RF pulse.

connections amongst ATCA and AMC subsystems must be increased to derive benefits from these standards. They make use of high speed serial interfaces available on the ATCA backplane that allow transferring data with throughput up to 10 Gbps.

The LLRF system for the XFEL should be designed to be maintainable and upgradeable for at least the next 20 years.

II. DATA TRANSMISSION CHANNELS

The FLASH accelerator works in a pulsed mode at a repetition rate between 1 Hz and 10 Hz, where the single RF pulse lasts approximately 2 ms [4]. Therefore, the designed LLRF control system requires three various types of communication links:

- real-time, intra-pulse links;
- real-time, inter-pulse links;
- non-real-time links.

This classification comes from different types of control algorithms implemented in the LLRF system. The relationship of real-time algorithms with a single RF pulse is depicted in Fig. 2.

The first category is real-time, intra-pulse links that are required for data transmission in the fast feedback loop or the interlock system. The latency for the intra-pulse links should be below 150 ns to provide required field stability and control bandwidth [11]. The second category is the real-time, inter-pulse links where data should be transmitted between two subsequent RF pulses. The third category includes non-real-time links which provide the data transport between slow parts of LLRF systems and other accelerator subsystems within a few RF-pulses.

Communication amongst the carrier boards in the LLRF system utilizes two categories of interfaces.

- Low latency communication that corresponds to intra-pulse links. The fast feedback algorithm executed by the LLRF system requires multiple real-time connections amongst the computation units. Low latency and guaranteed data arrival are required. The bandwidth of the link is estimated to be 3 Gbps with low duty cycle. During the inter-pulse dead-time, the links can be used for other purposes, but inter-feedback transmissions have the highest priority. For the ATCA standard such links can be implemented as full-mesh backplane links using user-defined lines. The connections must be implemented either as custom backplane or external inter-board links for the system implemented in the VME standard. A copper or optical fiber transmission medium can be used.

- Communication with control and diagnostic systems that covers inter-pulse and nonreal-time links. Data gathered by sensor units are transmitted to CPU blades for diagnostic, monitoring and presentation purposes. The control system processes the data and forwards results to other subsystems. The latency of the transmission is not important; however the data arrival time must be guaranteed. The link bandwidth estimated for the system, which is processing signals measured at the RF station consisting of 32 cavities, is 400 Mbps at a 10 Hz repetition rate and includes all measurement signals for individual cavities and important intermediate computation data. In the ATCA standard the links can be implemented as multiple PCIe or Gigabit Ethernet links which provide large bandwidth. The system based on the VME standard communicates with the VME master (for example the CPU in the VME crate) using a single parallel bus. The bus is time-multiplexed amongst the boards. The theoretical achievable bandwidth using VME in Direct Memory Access (DMA) mode is 320 Mbps, but practical measurements show, that only around 200 Mbps is reachable. Additional custom, external links are required when the full LLRF system is implemented using the VME standard. A copper or optical fiber transmission medium could be used, but additional copper or optical cables and connectors would decrease the reliability and availability of the whole system.

III. ATCA-BASED LLRF CONTROL SYSTEM ARCHITECTURE

The ATCA standard offers many important features that enhance the availability of the LLRF controller:

- modular design;
- supervision and monitoring of power supply;
- redundancy of important submodules and connections;
- built-in diagnostics.

Since a significant number of analog signals is connected to the LLRF controller, it is required to design a distributed controller composed of a few ATCA blades. ATCA carrier boards and AMC modules are used to maintain modularity and upgradeability [10], [12]. The complex architecture of the controller requires communication over a significant number of digital and analog links. The LLRF system is composed of:

- four custom-designed ATCA carrier boards (blades) with three AMC slots each;
- commercial Gigabit Ethernet switch in ATCA form with three AMC bays ¹ [13];
- ATCA processor blade with dual AMC bays and dual Xeon processor ² [14];
- off-the-shelf CPU board with PCIe Root Complex ³ [15].

The PICMG 3.0 specification defines the ATCA backplane that contains a number of general purpose serial links. The ATCA backplane is divided into three zones. Zone 1 is dedicated for a dual redundant power supply bus and redundant Intelligent Platform Management Interface (IPMI) bus for management and supervision. Zone 2 provides the data transport interfaces

¹Model ATS1936 from Diversified Technology.

²Model CPU-6900 from ADLink Technology.

³Model MPC8568E-MDS-PB from Freescale.

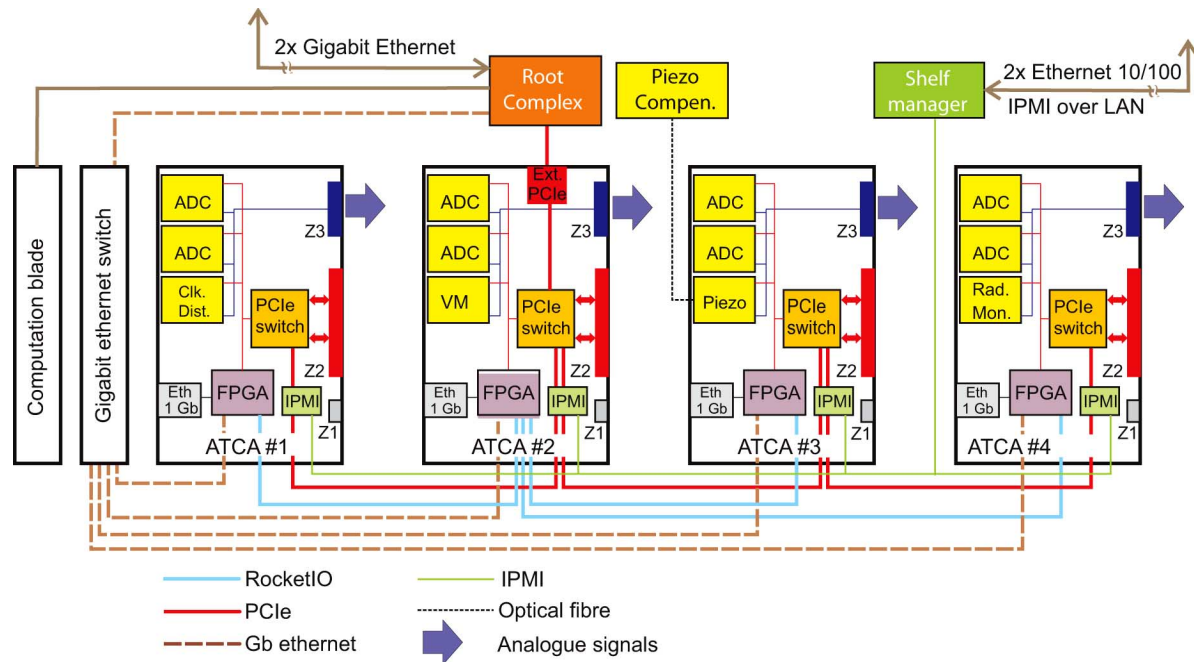


Fig. 3. A block diagram of the hardware of the ATCA-based LLRF system.

that support: base, fabric, update channel and synchronization clock interface. Zone 2 data transmission channels support Dual Star, Dual-Dual Star or Full Mesh topologies. The base channel interface, comprised of four signal pairs, is dedicated for the 10/100/1000BASE-T Ethernet interface. Subsidiary PICMG 3.x standards define communication protocols that can be implemented in the Fabric Interface, e.g., Gigabit Ethernet, InfiniBand, PCI express (PCIe) and StarFabric [10], [16]–[21]. They employ four differential lanes for a single communication channel. These interfaces allow an enormous data throughput (in the range of Gbps per serial link). The space available in Zone 3 is user defined and is usually used to connect the ATCA board to a Rear Transition Module (RTM). Zone 3 area can also hold a user-defined backplane to interconnect boards with the signals that are not defined in the ATCA specification.

The block diagram of the distributed controller showing connections amongst the modules is presented in Fig. 3. The analog part of the LLRF controller, i.e., downconverters and RF signals connectivity is installed on Rear Transition Module (RTM) boards [22]. Analogue signals (transmitted, forward and reflected powers) are connected to the RTM, downconverted and sent to the carrier board via the user-defined Z3 connector. The clock and trigger distribution network with frequencies $f = 81$ MHz and $f = 10$ Hz respectively is implemented on a custom-designed backplane in Zone 3.

The main advantage of the proposed solution is the elimination of signals on the front panel to simplify and improve the transmission of analog signals to the ATCA module. There is no need to reconnect dozens of cables when the ATCA controller is removed from the ATCA shelf.

The PCI express available on the Fabric Interface is used for data transmission amongst ATCA carrier boards and AMC modules. The PCIe interface is used for transmission of the control data. The interface uses a star topology with a single

Root Complex (RC). The PCIe standard requires the RC for the bus management and configuration. ATCA blades communicate with other ATCA blades installed in the same ATCA shelf using available PCIe switches. Certainly, an appropriate backplane of the ATCA shelf must support the full mesh topology on the Fabric Interface. All PCIe devices are available in the same memory space. The PCIe standard uses point-to-point connections; therefore a PCIe switch is required when multiple endpoints are connected. An eight-port, 32-lanes PCI express switch⁴, is installed on each ATCA carrier board to connect three AMC modules; FPGAs (Field Programmable Gate Array) with the Fabric connector (two channels PCIe x1) [23]. The PCIe x1 interface is available on Port 0 on the Fabric channel.

An additional PCIe x4 connector is reserved for the Root Complex required by the PCIe standard. The RC is connected to the ATCA carrier #2 using a dedicated PCIe cable [24]. However, the PCIe interface is not available outside the ATCA crate. Therefore, a software PCIe-to-Gigabit Ethernet bridge running on the RC board was designed to provide access to the PCIe bus for external applications. A hierarchical structure of layers required for the communication with the applications running on the external computation unit is presented in Fig. 4.

Algorithm Execution Details

A large number of the control algorithms running in the ATCA-based LLRF control system requires a significant computation power and several different communication links for one RF station. Due to large number of analog signals and different types of control algorithms it is not possible to implement the LLRF controller on a single ATCA carrier blade. In the present system all control algorithms are distributed among four carrier blades (ATCA #1–#4) and an optional CPU blade

⁴PEX8532, fabricated by PLX Company.

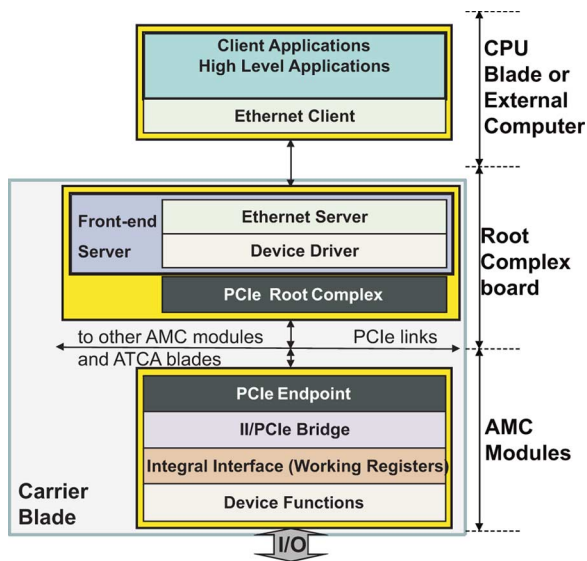


Fig. 4. A hierarchical structure of layers required for the communication with external applications.

located in the same ATCA shelf. Distributed algorithms require a number of communication links with different latencies and throughputs according to the defined classification.

All ATCA boards collect data from analog channels, calculate a partial vector sum and transfer the data to the main controller (ATCA #2). The controller, implemented in the FPGA device, receives the data using Low Latency Link and generates output signals that are connected to the Vector Modulator. Analogue signals connected to AMC modules with Analogue-to-Digital-Converters (ADCs) are digitized and sent to the FPGA using a differential signal. The partial vector sum calculated in the FPGA is transferred to the controller carrier blade ATCA #2. Therefore, a custom, Low Latency Link protocol is required for the transmission of the critical data, especially in the fast feedback loop, see Fig. 1. The LLL protocol required for the communication between the AMC modules and the carrier blades is provided by Fat Pipe's ports available on the AMC A+ type connector. The Full Mesh topology of the ATCA backplane supports LLLs with a custom protocol on the Fabric Interface required for the transmitting of low latency, real-time signals. The communication between carrier blades is achieved through Xilinx RocketIO ports on the Fabric Interface. For the intra-pulse operation a Full Mesh topology is involved. Intra-pulse communication channels do not require high throughput, but latency is more critical for these interfaces. In order to meet this highly demanding latency requirement the Xilinx Virtex 5 FPGA chip is used on the ATCA carrier blade.

During the RF pulse a large digital data set collected from several ADCs is stored in memory for further processing. The storage memory for the present system is available on AMC modules equipped with ADCs. After each pulse, stored data are transported from all buffers on AMC modules to FPGAs on ATCA carrier boards and the CPU blade. These data are used by different control algorithms implemented in FPGA chips and the CPU blade. Since the final allocation of the control algorithms

has not been decided yet, the computation load of particular carrier blades is also not yet known. It has been assumed that each carrier blade may require data from any other blade in the ATCA shelf. Since the latency of transferred data is not critical for inter-pulse links and may be higher than a few μs , the PCIe bus can be used. The throughput offered by a single lane PCIe interface is large enough for inter-pulse transmission. The measured latency for the PCIe connection is 1 μs during writing and 2 μs during reading (PCIe x1 transmission between the Xilinx V5 FPGA endpoint on the custom-designed AMC module and MPC8568E-MDS-PB) [15], [25]. PCI express has been chosen as a main protocol for the transmission of digital data between ATCA carrier blades and AMC modules within one ATCA shelf. The PCIe standard allows peer-to-peer connections. Therefore, a PCIe switch PEX 8532 is applied on each carrier board to realize a daisy chain topology on the Fabric Interface with latency in range of a few μs . The PCIe protocol is used for the real-time, inter-pulse data transmission such as control signals, data acquisition, diagnostics signals and piezo detuning calculation.

The last type of the non-real-time communication, with latencies in the range of milliseconds, is supported by Ethernet links in the ATCA shelf and the PCIe-to-Gigabit Ethernet bridge. The ATCA-based LLRF system contains an ATCA Ethernet switch located in the slot 1 of the ATCA shelf (a redundant switch may be installed in the slot 2). The PICMG 3.0 specification defines obligatory Ethernet links on the Base Interface located in Zone 2. The Base Interface allows an implementation of 1000BASE-T Ethernet links concentrated in the Ethernet switch. Physically, Ethernet links form a star topology with a central Ethernet switch. However, it is possible to achieve parallel communication between several slots in the ATCA shelf if the Ethernet switch implements non-blocking switching. Ethernet communication is dedicated for slow algorithms and the star topology on the Base Interface seems to be sufficient for currently planned control algorithms.

However, the current version of the ATCA carrier blade is not equipped with the Root Complex required for the PCIe interface; therefore an additional processor in the ATCA system is needed to serve the RC functionality. There are several options for RC implementation in the ATCA system. First, the RC may be installed on one of the commercially available AMC modules equipped with a processor with the PCIe interface. Such an AMC module should be installed on one of carrier blades in the AMC bay since each of the AMC bay is connected to the PCIe switch located on the blade. The disadvantage of this solution is that, the AMC module occupies one of the limited number of AMC bays and a failure of the particular carrier blade will disturb the PCIe communication in the entire ATCA shelf. Second, a commercial ATCA CPU blade which has PCIe links wired from the on-board CPU to the Fabric Interface may be used. In this case the configuration of the PCIe communication in the ATCA shelf is independent of the carrier blade. Unfortunately, there are a very limited number of commercial ATCA CPU blades which use PCIe links on the Fabric Interface [14]. Third, one can use a PCIe cable for the connection with an external processor working as the RC. This option is inconvenient due to the involvement of the external computer and potential

unreliability, however it is acceptable during the development of the system.

A. Protocol Details

The PCIe-to-Gigabit Ethernet bridge is a server application, working under the control of a Linux operating system, designed to perform the communication amongst external programs connected to the network and PCIe devices transparent to users. Several low-level drivers were designed to exchange data between the bridge application and hardware components connected to the PCIe bus. Drivers are responsible for data formatting, interrupt handling and reading from or writing to data memory locations. Appropriate procedures have been implemented for various types of devices. The communication amongst external client applications and the bridge server is performed using a High Level Application PCIe library (`libhlapcie`). The custom-developed library written in C++ programming language, is designed to perform data encapsulation, device or register address mapping and data format conversion. At the beginning of data transmission, the client application sets the address of the requested PCIe register and, if possible, establishes communication. The correct address contains three elements: the name of the ATCA carrier board, name of the requested device (e.g., a FPGA device on AMC or ATCA) and name of the register in the PCIe namespace. The communication library converts the name of the ATCA carrier board to the IP address and the name of the register to the offset relative to the PCIe base address. A transmission frame is sent to the bridge application. An additional communication protocol has been introduced to unify the data frame format and the addressing convention. The PCIe-to-Gigabit Ethernet bridge, based on the contents of the received frame, exchanges information with an appropriate device and sends answers back. The High Level Application library together with the bridge server application constitutes an intermediate level of the PCIe communication software subsystem, that can be used with applications like DOOCS servers, Matlab scripts or C/C++ standalone programs to visualize or control behaviour of hardware devices connected to the PCIe bus.

The transmission of the real-time data in the main controller loop requires a connection with latency below 150 ns, e.g., for the transmission of the partial vector sum from the Data Acquisition Boards (ATCA carriers #1, #3 and #4), see Fig. 3. Since none of interfaces offered by the ATCA standard is able to fulfil the requirement for Low Latency Links (LLL) for the main controller loop, a custom-developed protocol is used, based on Xilinx RocketIO transceivers available in the Virtex 5 family FPGAs [26]. The LLL protocol is implemented in a full-mesh topology on a fabric channel (port 1). The same protocol has been used for the data transmission from AMC modules to the data processing FPGA. All analog signals are delivered to AMC modules via the Zone 3 connector. Two additional cross-switches compatible with both RocketIO transceivers and the PCIe standard have been used to dynamically connect these signals to the required backplane transmission channels. The dynamic configuration allows ATCA carriers to be installed in any ATCA shelf slot.

The optional Gigabit Ethernet link shown by a dashed line in Fig. 3, is available on the Base Interface. All ATCA boards are connected to the ATS1936 Ethernet switch. The interface has no assigned functionality, it may be used for transmission of diagnostic or control data. A software or hardware processor with an Ethernet stack (i.e., TCP/IP) and a hardware Gigabit Ethernet controller implemented in the FPGA is required for data transmission. The processor with the Ethernet stack consumes a significant amount of the resources available in the Xilinx Virtex 5 chip family. The Xilinx V5 devices contain a hardware PCIe endpoint; therefore this protocol is more suitable for diagnostic purposes than the Gigabit Ethernet. There is no need to use software or hardware processors for the PCIe connection. Moreover, the latency of the PCIe link is lower and more predictable than for the Gigabit Ethernet.

The ATCA standard requires a redundant Shelf Manager (ShM), Intelligent Platform Management Controllers (IPMCs) to be installed on the ATCA carriers and Module Management Controllers (MMCs) on the AMC modules for management, supervision and monitoring of electronic components [27]. The ShM, IPMC and MMC are connected using the IPMI bus via a redundant System Management Bus (SMB). The ATCA standard offers a redundant connection for the most critical subsystems including the IPMI bus. The operator can monitor and control all subsystems via redundant 10/100 Mb Ethernet connections supported by IPMI over the LAN protocol, see Fig. 3. The ATCA standard also provides a redundant power supply bus to enhance the reliability of the system [10]. For example, a suitable warning is generated when the failure of one power supply unit is detected. Every ATCA carrier board is equipped with a diagnostic interface compatible with the EIA RS 232 standard.

The ATCA-based controller is connected to any other devices installed outside the ATCA shelf with optical fiber link. Fig. 3 shows such connection to Piezo Compensation subsystem.

B. FPGA Control Interfaces

Low-latency parts of the control algorithms are implemented in the FPGA fabric. An important feature of the design is the simple access to various registers located there, regardless of whether the external interface is VME, USB, PCI or PCIe. This design uses an *Integral Interface* (II) [28]. The Integral Interface is a set of VHDL functions and code generator programs, allowing easy access to a set of registers and memory blocks in the FPGA code. An interface to the user VHDL code is a simple bus interface, see Fig. 5. Internally, registers and memory blocks are visible as elements of two structures, s for scalar elements and v for memory blocks (vectors), so accessing them is straightforward. The set of registers and memory blocks is described by the user using the textual format exemplified in Fig. 6. This input file is subsequently processed by the code generator, which assigns addresses to registers, generates the VHDL template code and generates header files allowing access to registers from the C code.

As the user logic interface is not directly compatible with the PCIe bus, a special PCIe-to-Integral Interface bridge has been

```

iclk      : in  std_logic;
ii_resetN : in  std_logic;
ii_addr   : in  std_logic_vector(31 downto 0);
ii_writeN : in  std_logic;
ii_data_in : in  std_logic_vector(31 downto 0);
ii_data_out : out std_logic_vector(31 downto 0);
ii_strobeN : in  std_logic;
ii_ackN   : out std_logic;
ii_irqN   : out std_logic;
ii_irq_ackN : in std_logic;

```

Fig. 5. The Integral Interface user logic signals.

```

@version "1.0"
@mapFileName "DAC001_32.map"
reg id=reg1 width=14 attr=rw type=u impl=rw
comment="Register 1"
area id=areal width=12 attr=rw type=f impl=rw
count=234 addr=0x400 comment="Array 1"

```

Fig. 6. An example of definition of Integral Interface registers.

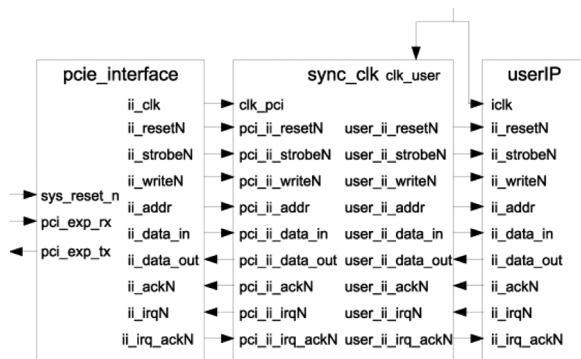


Fig. 7. The Integral Interface—the block diagram of the PCIe bridge system.

developed. The system uses the Xilinx Virtex 5 FPGAs, containing built-in hardware PCIe endpoint blocks, which handle PCIe configuration space requests, and allows communication with the PCIe bus via transaction layer packets. The developed bridge handles the conversion between the transaction layer and the user logic interface. The endpoint block transaction layer can work with a set of fixed frequencies —62.5 MHz, 125 MHz or 250 MHz. If the user logic requires an operation at a different frequency, e.g., enforced by the ADC sampling rate, an additional synchronizer block is required (see Fig. 7).

IV. EXPERIMENTAL RESULTS

Several experiments have been carried out to evaluate the performance of the proposed architecture. The application that has been tested using the ATCA architecture was the compensation of the Lorentz force detuning in 24 cavities simultaneously using the piezoelectric actuators mounted in the 8-cavity cryogenic accelerating module called ACC3, ACC5 and ACC6 of the FLASH accelerator. The pulse parameters have been determined by the operator; in the future this process will be implemented. The latency of the direct data transmission via the PCIe bus (e.g., the write operation) was measured to be $\sim 1 \mu\text{s}$. The latency is $\sim 2 \mu\text{s}$ for read operations. Therefore, to read large amounts of data from the algorithms implemented in the FPGA, a DMA transfer must be used. The measured throughput for the four-lane PCIe connection was 2400 Mbps in the DMA mode (the PCIe x4 transmission between the Xilinx

V5 FPGA endpoint on the custom-designed AMC module and the MPC8568E-MDS-PB computer) [15], [25]. The interrupt latency of the MP8568 PowerQUICC III processor working under the Montavista Linux operating system during the communication with the Virtex 5 FPGA connected via the PCIe interface turned out to be $\sim 1 \mu\text{s}$, sufficiently fast to implement a part of algorithms that have to operate between pulses, e.g., the computation of parameters for the Lorentz force detuning compensation on the PowerQUICC processor. The latency introduced by the PCIe switch PLX8532 is in range of 200 ns. According to the Xilinx datasheet the overall latency of the transmitter-receiver path for RocketIO blocks varies from 12.5 to 23 clock cycles (106.25 MHz) depending on the chosen configuration options.⁵ Therefore, the latency of the custom protocol is from 117 ns to 216 ns. The output update rate for the XFEL field controller is 200 ns. Assuming pipelined operation of the controller's computation core, the link fulfils timing requirements for the system.

V. CONCLUSION

Currently, the VME is the most popular system architecture used in the high energy physics' infrastructure. Despite simplicity and low cost, important features for demanding applications are missing (redundancy, fault tolerance, hot swap) and performance is rather low ($\sim 200 \text{ Mbps}$ typ.).

The ATCA architecture eliminates most of the VME weaknesses. ATCA offers much higher bandwidth realized using point-to-point links, therefore failure of one link does not cause the failure of the entire crate. Various standards of data transmission are available and it is possible to implement custom links. Apart from the front connector, it is possible to connect signals from the back of the crate, using the Rear Transition Module, which can help to eliminate unreliable front-panel connectors. ATCA uses a redundant power supply to enhance the reliability. One of essential features of the ATCA standard—the remote diagnostics via IPMI—allows early faults detection in this subsystem and in connection with the hot-swap feature, in many cases avoid the downtime entirely. ATCA has also its disadvantages. The barrier of entry is set higher, due to the price, which is higher today compared to the VME alternative, and its complexity. The redundant Shelf Manager was proposed to reduce the probability of system failure caused by damage of the IPMI subsystem.

Also, the standard written mainly for Telecommunication does not define analog signals. There is a work in progress on extensions of the standard for Physics—PICMG announces the formation of the xTCA for Physics Coordinating Committee [29].

The analysis of the PICMG 3.0 specification shows that a variety of obligatory and optionally defined communication links fulfil demanding requirements of control algorithms of the LLRF system. Obligatory links in the ATCA shelf provide the support for the redundant management and configuration subsystem. PCIe and Gigabit Ethernet interfaces are suitable for transferring inter-pulse or nonreal-time data. However, the application of PCIe or Gigabit Ethernet connections, even

⁵For device Xilinx XC5VLX50T-FFG665-3C with RocketIO clock equal to 106.25 MHz

though a high throughput, introduce a significant latency during the data transition. The latency in range of microseconds for PCIe or hundreds of microseconds for the Gigabit Ethernet is too much for real-time data transfer in the fast feedback loop of the LLRF controller. The latency should not exceed 200 ns. Moreover, an application of hierarchical PCIe or Ethernet switches significantly increases the latency.

The abundance of point-to-point links available on the ATCA backplane allows implementation of custom-defined latency-optimized protocols. The custom protocol is based on RocketIO transmitters offered by Xilinx. The measured latency fulfills requirements of the LLRF system. The latency could be improved when the RocketIO frequency will be increased. Xilinx allows to use the frequency up to 156 MHz [26], [30].

In the LLRF system almost 200 analog and digital cables will be connected to a single RF station. Therefore, the exchange of LLRF control board can be difficult, especially in an emergency situation. The application of the ATCA shelf with RTM modules and rear connections should simplify this procedure considerably.

The LLRF control system of XFEL accelerator must be maintained over the next 20 years. ATCA appears to be the most flexible platform available to meet upgrade requirements over this period.

Having in mind the above considerations, the prototype XFEL LLRF feedback controls system has been successfully implemented on an ATCA platform through custom designed AMC carriers containing high speed COTS ADC's and an associated RTM module that supports RF-IF down-conversion and transmission of analog sensor signals. The overall performance goals have been realized and compare favorably with an earlier system based on VME.

The main advantages of the high availability ATCA platform have been realized with the implementation of full mesh backplane, redundant power supplies, and IPMI management layer. Some testing of the latter still needs to be done. The one non-redundant controller is the Root Complex controller for PCIe, which will be eliminated in the next iteration of design. The overall architecture follows an earlier system except for details specific to the ATCA implementation. Three levels of protocols have been successfully implemented and tested: RocketIO, for the communication at the lowest latency level, PCIe for DAQ, and Ethernet for client applications. The controls algorithms were successfully implemented to perform the fast intra-pulse corrections and feed-forward operations to the compensating actuators.

There is still an open question whether the forty-year old VME standard will be supported within the next 20 years. ATCA and AMC standards are becoming more and more common. Currently more than 100 companies are participating in the development of these standards. Therefore, standards should be available within the next dozens of years. Finally, we can conclude that the ATCA standard appears to be more suitable to design a complex and the reliable LLRF control system for the XFEL accelerator.

REFERENCES

- [1] S. Simrock, "State of the art in RF control," in *Proc. LINAC 2004*, Lubeck, Germany, Aug. 2004.
- [2] S. Simrock, W. Cichalewski, M. Grecki, G. Jablonski, and W. Jalmuzna, "Universal controller for digital RF control," in *Proc. 8th Eur. Particle Accelerator Conf., EPAC*, Edinburgh, U.K., 2006, pp. 1459–1461.
- [3] R. Pietrasik, W. Giergusiewicz, W. Jalmuzna, K. Pozniak, R. Romaniuk, and S. Simrock, "Measurements of SimCon 3.1 LLRF control signal processing quality for VUV free-electron laser FLASH," *Proc. SPIE: Photonics Applications in Astronomy, Communications, Research and High Energy Physics Experiments*, vol. 6347, pp. 1–16, Jun. 2006.
- [4] W. Giergusiewicz, W. Jalmuzna, K. Poźniak, N. Ignashin, M. Grecki, D. Makowski, T. Jeżyński, K. Perkuszewski, K. Czuba, S. Simrock, and R. Romaniuk, "Low latency control board for LLRF system—SIMCON 3.1," *Proc. SPIE, Photonics Applications in Industry and Research IV*, vol. 5948, pp. 710–715, 2005.
- [5] S. Simrock, "Measurements for low level RF control systems," *Measurement Science and Technology on Metrological Aspects of Accelerator Technology and High Energy Physics Experiments, Special Issue*, pp. 2320–2327, Dec. 2006.
- [6] W. D. Peterson, *The VMEbus Handbook*. Scottsdale, AZ: VFEA Int. Trade Assoc., 1993.
- [7] S. N. Simrock, "Low level radio frequency control system for the european XFEL," in *Proc. 13th Int. Conf. Mixed Design Integr. Circuits Syst., MIXDES 2006*, Jun. 2006, pp. 79–84.
- [8] S. Simrock, M. Grecki, W. Jalmuzna, T. Jezynski, W. Koprek, and P. Pucyk, "Distributed versus centralized ATCA computing power," in *Proc. 15th IEEE-NPSS Real-Time Conf., 2007*, Batavia, IL, Apr./May 2007, pp. 1–6.
- [9] S. Simrock, "Requirements for the ATCA based LLRF evaluation system." Review of LLRF System Based on ATCA Standard, Nov. 8–9, 2007.
- [10] "AdvancedTCA Base Specification, PICMG 3.0," PICMG, Jan. 2003.
- [11] E. Vogel, "High gain proportional RF control stability at tesla cavities," *Phys. Rev. ST Accel. Beams*, vol. 10, no. 5, pp. 052001–0520012, May 2007.
- [12] "AdvancedMC Mezzanine Modul Specification, PICMG AMC.0," PICMG, Nov. 2006.
- [13] "ATS1936 ATCA Board 10 Gigabit Switch—User's Guide," Diversified Technology, Inc., 2007 [Online]. Available: <http://www.dtim.com/products/atca/ats1936.php>
- [14] "ATCA-6900 Series, Dual Quad Core LV-Xeon AdvancedTCA Processor Blade With AMC Bays—User's manual," ADLINK Technology Inc., 2008 [Online]. Available: <http://www.adlinktech.com>
- [15] "PowerQUICC MDS Platform I/O Board—User Manual," Freescale Semiconductor, 2005 [Online]. Available: <http://www.freescale.com>
- [16] "AdvancedTCA Base Specification, PICMG 3.0," PICMG, Dec. 2002.
- [17] "AdvancedTCA Ethernet Specification, PICMG 3.1," PICMG, Jan. 2003.
- [18] "AdvancedTCA InfiniBand Specification, PICMG 3.2," PICMG, May 2003.
- [19] "AdvancedTCA StarFabric Specification, PICMG 3.3," PICMG, May 2003.
- [20] "AdvancedTCA PCI Express Specification, PICMG 3.4," PICMG, Jan. 2003.
- [21] "AdvancedTCA RapidIO Specification, PICMG 3.5," PICMG, Sep. 2005.
- [22] "AdvancedMC Rear Transition Module Specification, PICMG ARTM.0," PICMG, Jul. 2009.
- [23] "ExpressLane PEX8532AA/BA/BB/BC 8-Port/32-Lane Versatile PCI Express Switch—Data Book," PLX Technology, Inc., 2007 [Online]. Available: <http://www.plxtech.com/products/expresslane/PEX8532.asp>
- [24] "PCI Express External Cabling 1.0 Specification," PICMG, Feb. 2007.
- [25] D. Makowski, A. Piotrowski, and A. Napieralski, "Universal communication module based on AMC standard," in *Proc. 15th Int. Conf. Design Integr. Circuits Syst., MIXDES 2008*, Jun. 2008, pp. 139–143.
- [26] "Virtex-5 FPGA RocketIO GTP Transceiver—User Guide," Xilinx, Dec. 2008.
- [27] "Intelligent Platform Management Interface Specification," PICMG, Feb., 2004.
- [28] A. Piotrowski, S. Tarnowski, G. Jablonski, and A. Napieralski, "Integral interface—universal communication interface for FPGA-based projects," in *Proc. 14th Int. Conf. Design Integr. Circuits Syst., MIXDES 2007*, Jun. 2007, pp. 115–119.
- [29] "XTCA for Physics," PICMG, Mar. 2009.
- [30] "Virtex-5 Data Sheet: DC and Switching Characteristics," Xilinx, Jan. 2007.