

A Review of Technologies for the Transport of Digital Data in Recent Physics Experiments

Denis Calvet

Abstract—This paper describes past and present techniques and technologies for the transport of digital data in recent physics experiments. After an overview of the typical requirements for modern data acquisition systems in the field of large scale experimental physics, we detail the successes and failures observed over the last 20 years of evolution of high-speed point-to-point link technology, networking standards and products. Modern data transport technology is presented along with several applications to experiments under construction. Advanced techniques, emerging technologies and trends in the field of high-speed digital data transport are outlined in the perspective of future experiments.

I. INTRODUCTION

THE basic principles for the design of data acquisition systems (DAQ) in experimental physics have not changed very much over the last three decades. The art of the DAQ designer is to cascade in a clever way behind a detector, electronic devices that fall in one of three categories: data processing devices, data retention devices and data transport devices. Data processing devices, such as discrete transistors, logic gates or computers, transform, and generally reduce, data to extract content information. Data retention devices, such as capacitors or digital memories, temporarily hold data that cannot be processed or transported as such, for example bursty data, or data awaiting the decision to be discarded or kept. Data transport devices, from simple wires to local area network links, are used to move/gather/route data to an environment appropriate for temporary storage or processing. This paper addresses the techniques and technologies used for the transport of digital data within large-scale physics experiments. What solutions were deployed in the past, and what performance was achieved? Which products and technologies were successful and for what reasons some others failed? What are today's key standards, products and techniques? How these are going to evolve? Without being an exhaustive survey of the abundant literature on the subject, this paper gives pieces of answers drawn from personal observations, readings, discussions and the general experience of the author.

II. REQUIREMENTS FOR THE TRANSPORT OF DIGITAL DATA IN PHYSICS EXPERIMENTS

Some of the relevant items to classify digital data transport devices include:

- Bandwidth: kbps, Mbps or Gbps-class
- Connectivity: point-to-point, multicast/broadcast, random traffic, fixed pattern,
- Latency: deterministic or variable,
- Medium: electrical, optical, radio,
- Distance: short (< 10 m), medium, large (> 2 km).

Among the 216 possible combinations of the previous parameters, it is both a pleasant and difficult exercise to spot any sensible configuration that was never exploited by at least one experiment in physics. Try! To remain general, this paper focuses on some of the most common configurations.

In modern large experiments, the transport of data off-detector typically uses tens to thousands of Gigabit-per-second-class point-to-point links running in parallel [1]. Optical fibers are often preferred over copper for compactness and to provide galvanic isolation. Timing distribution systems are an example of latency critical multicast/broadcast networks. These routinely fanout synchronous signals at a few 10 MHz rate from a central point to hundreds of units located several tens of meters apart, with nanosecond-order system-wide skew and sub-nanosecond jitter [2]. Trigger systems and event builders require data transport networks with multi-gigabit per second aggregate bandwidth, and tens to hundreds of ports [3].

III. UNTIL THE MID-80'S: THE BUS ERA

From the 50's to the mid-80's, the typical size of electronic systems for experimental physics has grown from one, to a few, then tens of crates. Backplane bus standardization was a key activity and interconnecting multiple crates with point-to-point links was the way to grow beyond the limitations imposed by a single enclosure. Several key standards were established by the nuclear science community: NIM (Nuclear Instrumentation Methods, standard DOE/ER-0457 established in 1964), CAMAC (Computer Automated Measurement And Control, European ESONE, US NIM standard EUR 4100/IEC 516 established in 1968) and FastBus (ANSI/IEEE 960-1986, IEC 935) [4], [5]. FastBus was successfully used in many major

Manuscript received June 3, 2005.

D. Calvet is with DSM/DAPNIA, Commissariat à l'Énergie Atomique / Saclay, F-91191 Gif-sur-Yvette, France (telephone: +33169086909, e-mail: calvet@hep.saclay.cea.fr).

physics experiments in the 80's (over 20 publications related to Fastbus appear in [6]), but the complexity of the standard and the lack of acceptance by the industry mass market made users' interest shift towards the VMEbus (Versa Module Europe) [7]. For these various bus standards, inter-crate links running at up to ~20 MB/s over several meters were based on parallel flat ribbon cables using TTL or ECL signaling levels. To interconnect VME crates, the VICbus [8] was popular. Up to 33 MB/s transfer rates could be achieved over several meters using 64 twisted-pair cables.

Commercial networking technologies included Fiber Distributed Data Interface (FDDI) [9] running at 100 Mbps over optical media, Ethernet based on a shared media (10 Mbps) [10], and Token Ring (4.16 Mbps), a technology originating from IBM standardized with minor variations by the IEEE 802.5 working group [11].

Optical links running at 100 Mbps were the highest practical speed available at relatively low cost thanks to the support of industry to the FDDI market. Dedicated chipsets such as Advanced Micro Devices TAXIchip™ [12] provided a ready-to-use solution for point-to-point links running at up to 140 Mbps over fiber-optic, coaxial cable or twisted pair media. Some applications of these devices are reported in [13], and notably in [14] where 128 optical links running in parallel were used for the readout of a time projection chamber.

IV. LATE 80'S TO MID 90'S: THE TRANSPUTER AND DSP ERA

Backplane buses played a central role for gathering detector data until bandwidth limitations became a critical bottleneck. To increase throughput, switched-based systems with multiple high speed links were introduced. Standard buses were nonetheless kept for mechanics, power, cooling, configuration, slow control, monitoring, and less demanding DAQ tasks.

A. Transputers

One of the most innovative products in the history of micro-processors is certainly the Transputer (*trans-istor com-puter*) introduced by INMOS in 1984 [15]. The basic idea is to integrate on the same silicon die a microprocessor core, a memory and communication links. By assembling many of these individually simple devices, parallel systems capable of performing complex tasks could be devised. A Transputer had 4 bi-directional links, running concurrently with the CPU. Operating speed was 5 Mbps for the first generation of Transputers, and reached 100 Mbps for the last model (T9000). Cross-point switches were also available to built complex network topologies. The C104 chip integrated 32 ports at 100 Mbps on a single die [16]. Transputers and their companion switch devices found numerous applications in instrumentation. A few examples are the event builder of the Zeus experiment [17], the DAQ system of the GA.SP experiment [18], the second level trigger of L3 experiment [19], and on-line filtering in CPLEAR [20]. A 1024-node system based on the C104 packet switch was constructed [21].

Transputer links used a clever scheme for data encoding, called Data Strobe (DS) encoding: data is sent on the Data line unmodified; the Strobe line changes state only if a data bit has the same value as the previous one. The exclusive OR of the Data and Strobe lines provides the receiver with the reconstructed dual-edge clock for sampling the Data line. This scheme is simple to implement (no PLL), allows for a full bit period of skew tolerance, and is auto-baud rate for the receiver. Several successor standards exploit this data encoding scheme: IEEE Std. 1355 [22], Firewire [23], and SpaceWire [24].

The enthusiasm for Transputers lasted about 6 years. Excessive delays in the production of the T9000 (bugs in the silicon, slower clock rate than initially planned), and the introduction of fast Digital Signal Processors (DSPs) put an abrupt end to the concept.

B. DSPs for parallel systems

At least two very successful DSP devices for parallel systems are worth mentioning: Texas Instruments' TMS320C40 [25] introduced in 1991, and its aggressive competitor, Analog Devices' Sharc (ADSP-2106x), introduced in 1994 [26]. Both devices integrate on a single chip a 32-bit floating point processor core, a static RAM, and 6 high speed (30-40 MB/s) communication links. Commercial multi-DSP boards equipped with 4-10 DSPs became widely used. Some applications in physics are reported in [27] and notably in [28], where ~140 boards, each carrying 6 Sharc DSPs, were used to buffer and switch detector data in the Hera-B experiment.

Each C'40 link was a half-duplex port (8 data + 4 control lines) and could run at up to 30 MB/s. Each Sharc link was programmable in either transmit or receive mode, used 4 lines for data, 1 line for clock (up to twice the CPU clock rate of 40 MHz), and 1 acknowledge line. Unipolar TTL level signaling was used. Transfer rates of up to 40 MB/s could be achieved over few tens of centimeters using high quality ribbon cables.

The competitive advantage of using DSPs in instrumental physics lasted about half a decade. During several years, market demands drove the evolution of DSPs more towards lower power consumption and low cost than ultimate computing power and I/O bandwidth. Steadily increasing CPU power and the introduction of the PCI bus in 1992 brought decent I/O capabilities to PCs for a bargain price compared to multi-DSP boards. Today, DSPs running at GHz clock rates have brought back these devices on the forefront of the scene.

V. MID/LATE 90'S: THE ATM, SCI, FAST ETHERNET DEBATE

By the late 80's, there was a crucial need for a new generation of standards for high speed data transfers in many different sectors of information technology. Surprisingly not, each community came up with at least two different and competing standards. Vast R&D programs to evaluate some of these technologies for applications at the future Large Hadron Collider (LHC) were initiated by CERN [29], [30]. What happened to these (too numerous) products and standards?

A. Beyond the bus concept

Looking for a successor to Fastbus and Futurebus, a community, partly close to high energy physics, developed the Scalable Coherent Interface (SCI) [31], while a competing group proposed Quick Ring. SCI is based on 1 GByte/s links (originally ECL signaling) interconnecting devices in a ring topology, with possible bridging between rings. It has some support for cache coherency, features high bandwidth and extremely low-latency transfers. This makes SCI ideal for building multi-processor systems. Quick Ring is a slower and simpler version (no cache coherency). Unfortunately, both SCI and Quick Ring lacked the necessary support from industry and applications for the mass market to really take off. The main tribute to SCI is an extension of the initial standard that defines one of the two standards for Low Voltage Differential Signaling (LVDS) [32], [33]. To the author, the LVDS specification is undoubtedly the most influential document of the decade in the field.

B. Linking storage devices

Initially meant for supercomputer to mass storage transfers, research on gigabit technology at Los Alamos National Laboratory led to the High Performance Parallel Interface (HIPPI) standard [34]. HIPPI initially offered 800 Mbit/s of bandwidth over a 32 bit interface. Also meant to connect computers to storage devices, a consortium of industrial developed Fibre Channel [35]. The initial version ran at 266 Mbps and was followed by a 1 Gbps version. Although both standards still co-exist and are being developed, Fibre Channel is having more commercial success than HIPPI.

In DAQ systems for experimental physics, the chipsets and components developed for Fibre Channel / serial HIPPI found many applications. Examples of devices for high speed point-to-point links are Cypress Hotlink (266-400 Mbps) and the extremely popular Hewlett Packard's G-link (1 Gbps). In order to avoid that the obsolescence of a device renders the design of a board unusable and to benefit easily from new faster devices, the S-Link concept was proposed [36]. It defines a FIFO-like interface which is independent of the physical link layer. Board designers just need to place connectors on their motherboard following the specification. Then any S-Link compliant third-party mezzanine card that includes the physical-layer-of-the-year can be used. The concept found many applications (e.g. in [37]), and numerous products are still available.

C. Unifying wide area / local area networking (WAN/LAN)

Promising the convergence of voice, data and video traffic over a common media, a forum (with strong representatives from the telecom world) established Asynchronous Transfer Mode (ATM) technology [38]. In ATM, data is carried by short (53 bytes) cells, time-multiplexed over the same media. Quality of service mechanisms allow to transfer multi-media, voice and data traffic across a unified infrastructure. Commercial switching products had up to 256 bi-directional

155 Mbps ports. Despite massive support from the telecom industry and some LAN equipment vendors, ATM failed to capture more than a small fraction of the LAN market. Nonetheless, ATM was the technology of choice for backbone infrastructure (e.g. Internet Service Providers) until the late 90's, and even today, the ATM service market is still a very profitable business. In physics applications, ATM products have been used for event builders in [39] and [40].

D. Upgrading proven LAN technology

Based on the success of 10 Mbps Ethernet, the local area network community proposed a ten-fold increase of performance with 2 competing standards: Fast Ethernet [41] and 100VG-AnyLan. The advantage of the 100VG-AnyLan proposal was to support both Ethernet and Token Ring frame types, but this argument did not appeal much to customers. Fast Ethernet supports half-duplex and full-duplex operations and is interoperable with the first generation of Ethernet (identical frame format, auto-negotiation of speed on each segment). Another major evolution was the introduction of switched Fast Ethernet. Instead of sharing bandwidth between devices attached to the same segment, star topologies based on multi-port switches could be built. A modest size 16-port Fast Ethernet switched network provides a 320-fold increase of bandwidth compared to a single segment half-duplex 10 Mbps Ethernet! Fast Ethernet captured the largest fraction of the LAN market very rapidly, reducing month after month the chances of success of ATM. Fast Ethernet technology is ubiquitous in today's networking and one would hardly find systems for experimental physics where no single bit of data is sooner or later carried over a Fast Ethernet connection.

E. Proprietary products

Adding to the profusion of new standards that appeared in the 90's, several companies put their own products in the arena: Mercury's Raceway (ANSI/VITA Standard 5-1994) used in [42], Myricom's Myrinet (ANSI/VITA Standard 26-1998) used in [43], Sky Computers' SkyChannel Packet Bus (ANSI/VITA Standard 10-1995), etc. Despite the merit of each product, deployment in physics experiments was marginal, and choosing exotic technologies often proved (not always...) to be a good recipe for having regular system "upgrades"!

VI. LATE 90'S TO PRESENT: THE (MULTI-)GIGABIT ERA

A. LVDS technology and FPGA's

As mentioned earlier, the LVDS specification had a profound impact on a whole sector of the silicon industry. Families of devices to transport data at (multi-)Gbps rate over up to several meters of copper cable were introduced by National Semiconductor [44], and lower speed devices are available from Texas Instruments. These devices are cheap, flexible and draw very low power. An example of application is reported in [45] where a system based on 320 LVDS

serializers and de-serializers handles 300 Gbit/s of data with an input-to-output delay of 200 ns. LVDS also supports multi-drop applications and high speed backplane bus applications (up to 5 Gbps per bus line pair).

The progresses made on Field Programmable Gate Arrays (FPGA's) are also spectacular. Two major evolutions took place recently: the ability of FPGA's to support LVDS and other high speed signaling I/O standards, and the integration of dedicated blocks to complement programmable logic: RAMs, multi-gigabit class transceivers, DSP slices/RISC processors; a concept referred to as "a system-on-a-chip". Multi-million gate FPGA devices capable of digesting several 10 Gbps of I/O bandwidth offer ultimate flexibility for a few hundred dollars. FPGA's are an indispensable ingredient in countless industrial applications, consumer products, and modern DAQ systems in experimental physics.

B. Gigabit Ethernet

Following its predecessors, the Gigabit Ethernet standard introduced in 1998 had a large and immediate commercial success. In less than a year, Gigabit Ethernet put a definitive end to the market of ATM at the core of enterprise networks. The reasons for the success of Gigabit Ethernet are manifold: the technology did not ambitioned to be universal; the standard converged very quickly; the technology is rather simple, bears a famous name, and is compatible (at the frame level at least) with previous generations of Ethernet. The original CSMA/CD scheme was reworked for Gigabit Ethernet, but the real intended use of the technology is full duplex point-to-point links interconnecting hosts via switches. Like Fibre Channel, 8B/10B encoding is used, but while the bandwidth of Fibre Channel links is given after encoding, Gigabit Ethernet actually transfers 1,000,000 bits of data per second in both transmit and receive (i.e. the actual line rate is 1.25 Gbaud each way). Several types of media are standardized: short and long wavelength optical fibers, short run copper (<25 m) and 4-pair category 5 unshielded twisted pair (up to 100 m).

The DAQ of the Compass experiment uses 4 16-port Gigabit Ethernet switches [37]. Most of the large scale experiments under construction plan to use Gigabit Ethernet in their DAQ system; some R&D work is presented in [3]. High-end commercial switches now offer up to 400 Gigabit Ethernet ports and 700 Gbps of aggregate backplane bandwidth.

C. Optical technology

Many progress on optical interconnects were also made during the last decade. Public networks built on Synchronous Optical Network - Synchronous Digital Hierarchy (SONET-SDH) standards were developed [46]. The OC-3 rate (155.52 Mbps) popular in the mid 90's was quickly followed by OC-12 (622.08 Mbps), then OC-48 (2.48 Gbps), OC-192 (10 Gbps), and products are now available for OC-768 (40 Gbps).

Dense Wavelength Division Multiplexing (DWDM) is a technique to increase the throughput of optical links in a

scalable way by combining many wavelengths (up to 160) onto a single fiber. Products are almost exclusively used for long haul transport. The technique is used to transport the data of the Antares underwater neutrino telescope located 40 km off-shore (6 wavelengths per fiber) [47], and an evaluation of an OC-48 DWDM transponder in view of future DAQ systems is given in [48]. For less demanding applications, full-duplex transceivers using a single fiber are available from several vendors in speeds ranging from 100 Mbps to 1.25 Gbps.

Parallel optics is also taking off. The SNAP12 Multi Source Agreement (MSA) is a specification followed by many vendors for interoperable parallel optics based on 12-fibers composite cables (MTP®/MPO). Current products include 12-channel transmitters and receivers (up to 3.125 Gbps per channel) and transceivers with 4 duplex channels. Transceivers with up to 36 duplex 2.7 Gbps channels are also being introduced by Tyco/AMP. An evaluation of a 22-bit optical transceiver was reported in [49]. An application of parallel optical transceivers and passive optical cross-connects is presented in [50].

VII. EMERGING TECHNOLOGIES AND STANDARDS

Recent standards for high speed digital electronics deal with signals above 300 MHz, and data rates in the 1-40 Gbps range. The general trend is high speed serial I/O. For mass market products, Firewire will soon deliver 1.6 Gbps and possibly 3.2 Gbps. USB2 (currently at 480 Mbps) may also evolve, although wireless USB could become the dominant trend. For interfaces to mass storage devices, Serial ATA is evolving from today's 1.5 Gbps rate towards 3 Gbps, while 4-10 Gbps Fibre Channel products are expected this year. The standard for 10 Gbps Ethernet was ratified in 2002, and products are now reaching maturity. The evolutions of the PCI bus are also promising: 3 different directions are being pursued (see www.pcisig.com). "Conventional PCI" is the evolution of the original specifications; PCI-X is a backward compatible version, operating at 133 MHz, 266 MHz or 533 MHz. Because the limits of the parallel bus concept may soon be reached, a radical change is proposed with PCI-Express™: the multi-drop parallel bus paradigm is abandoned in favor of multiple serial lanes (typical 2.5 Gbps per lane) transporting data packets. A competitive, and possibly complementary, standard is RapidIO® (see www.rapidio.org). Quality of service and multi-cast are among the new features offered. The InfiniBand™ architecture is heavily promoted (see www.infinibandta.org). Specifications for interoperable parallel copper cable and parallel optics modules are being set up by the IBPACK Multi Source Agreement group (see www.ibpak.org). Many of the new optical networking standards are being set up by the Optical Interconnecting Forum (OIF, see www.oiforum.com).

VIII. SUMMARY AND OUTLOOK

Over the last 25 years, the I/O bandwidth handled by electronic components and systems has increased by 2 to 3 orders of magnitude. Have the limits of parallel buses finally been reached? Will PCI-Express™ be more successful than SCI? Will DSPs with RapidIO® links bring again DSPs on the forefront of I/O intensive real-time computing? Or will FPGA's keep their current leadership? How fast and far can copper really go? Will parallel optics develop? Success is unpredictable, but experimental physics will, for sure, continue to benefit from the latest technological advances in the field.

IX. REFERENCES

- [1] P. Moreira et al., "G-Link and gigabit Ethernet compliant Serializer for LHC Data Transmission", IEEE Nucl. Sci. Symp. conference record, vol. 2, pp. 9/6-9/9, October 2000.
- [2] B.G. Taylor, "TTC Distribution for LHC Detectors", IEEE Trans. Nucl. Sci., vol. 45, no. 3, pp. 821-828, June 1998.
- [3] A. Barczyk et al., "The New LHCb Trigger and DAQ Strategy: a System Architecture Based on Gigabit Ethernet", IEEE Trans. Nucl. Sci. vol. 51, no. 3, pp. 456-460, June 2004.
- [4] P. Pointing and H. Verweij, "Instrumentation Buses for High Energy Physics, Past, Present and Future", IEEE Trans. Nucl. Sci., vol. 38, no. 2, pp. 322-324, April 1991.
- [5] R.S. Larsen, "Fastbus – Status of Development of a standard High Speed Data Acquisition Bus for High Energy Physics", IEEE Trans. Nucl. Sci., vol 25, no. 1, pp. 735-, Feb. 1978.
- [6] IEEE Trans. Nucl. Sci. vol. NS33 No. 1, February 1986.
- [7] "The VME bus specification", ANSI/IEEE 1014-1987, IEC 821.
- [8] C. F. Parkman, "VICbus: VME Inter-Crate Bus, A Versatile Cable Bus", IEEE Trans. Nucl. Sci. vol. 39, no. 2, pp. 77-84, April 1992.
- [9] R. Jain, "FDDI Handbook: High Speed Networking with Fiber and Other Media", Addison Wesley, Reading MA, April 1994.
- [10] "Carrier Sense Multiple Access with Collision Detection (CSMA/CD) method and physical layer specifications", ANSI/IEEE Std. 802.3, first approved 1983, current version 8 March 2002. Available online: <http://standards.ieee.org/getieee802/index.html>
- [11] "Token Ring Access Method and Physical Layer Specifications", ANSI/IEEE Std. 802.5, first approved 1985, current version 26 May 1998. Available online: <http://standards.ieee.org/getieee802/index.html>
- [12] "TAXIchip™ Integrated Circuits, Transparent Asynchronous Transmitter/Receiver Interface, AM 7968/ AM 7869 Datasheet and Technical Manual", Advanced Micro Devices, 1994.
- [13] C. Swoboda et al., "FastBus Standard Fiberoptic Data Transfer Link", IEEE Trans. Nucl. Sci., vol. 36, no. 1, pp. 774-779, February 1989.
- [14] F. Bieser, R. Jones and C. McParland, "Data Links for the EOS TPC", IEEE Trans. Nucl. Sci. vol. 38, no. 2, pp. 335-336, April 1991.
- [15] R. Yvemey-Cook, "Legacy of the Transputer", Architectures Languages and Techniques, volume 57 on concurrent systems engineering techniques, B.M. Cook, IOS Press, 1999.
- [16] M.D. May et al., "Networks, Routers and Transputers", IOS Press, Amsterdam, 1994.
- [17] U. Behrens, L. Hagge et W.O. Vogel, "The Zeus Event Builder: Experience with a Distributed Real-time Parallel Transputer System", IEEE Trans. Nucl. Sci., vol. 41, no. 1, pp. 239-245, February 1994.
- [18] D. Colombo et al., "The Transputer based GA.SP Data Acquisition System", IEEE Trans. Nucl. Sci. vol. 39, no. 2, pp. 103-108, April 1992.
- [19] J.J. Blaising et al., "Performance of the L3 2nd level trigger implemented for the LEP II with the SGS Thomson C104 packet switch", IEEE Trans. Nucl. Sci., vol. 45, no. 4, pp. 1765-1770, August 1998.
- [20] R. Heely et al., "The Application of the T9000 Transputer to the CPLEAR experiment at CERN", Nucl. Inst. and Meth. In Physics Research, vol. A 368, pp. 666-674, 1996.
- [21] M. Zhu et al., "Realisation and performance of IEEE 1355 DS/HS link based, high speed, low speed, low latency packet switching networks", IEEE Trans. Nucl. Sci. vol. 45, no. 4, pp. 1849-1853, August 1998.
- [22] "Standard for Heterogeneous InterConnect (HIC), Low cost low latency scalable serial Interconnect for parallel system construction", IEEE Std. 1355-1995.
- [23] "Standard for a high performance serial bus – Firewire", IEEE Std. 1394-1995.
- [24] "SpaceWire – Links, nodes, routers and networks", Std. ECSS-E-50-12A, European Cooperation for Space Standardization, 24 January 2003. Available online: <http://www.ecss.nl>
- [25] "Parallel Processing with the TMS320C4x", Application Guide, Texas Instruments, February 1994.
- [26] "Considerations for selecting a DSP Processor – Why buy the ADSP-21060, Analog Devices ADSP-21060 Sharc versus Texas Instruments TMS320C40", Application note AN-403, Analog Devices, August 1995.
- [27] K. Blacker et al., "The Jet Fast Central Acquisition and Trigger System", IEEE Trans. Nucl. Sci. vol. 41, no. 1, pp. 111-116, February 1994.
- [28] V. Egotychev et al., "Architecture of the Hera-B DAQ System", IEEE Trans. Nucl. Sci. vol. 50, no. 4, pp. 859-863, August 2003.
- [29] J. A. C. Bogaerts et al., "SCI based Data Acquisition Architectures", IEEE Trans. Nucl. Sci. vol. 39, no. 2, pp. 85-94, April 1992.
- [30] M. Letheren et al., "An Asynchronous Data-Driven Event-Building Scheme based on ATM Switching Fabrics", IEEE Trans. Nucl. Sci. vol. 41, no. 1, pp. 257-266, February 1994.
- [31] D. B. Gustavson, "IEEE P1596, a Scalable Coherent Interface for GigaByte/sec Multiprocessor Applications", IEEE Trans. Nucl. Sci. vol. 36, no. 1, pp. 811-813, February 1989.
- [32] "IEEE Standard for Low Voltage Differential Signalling (LVDS) for Scalable Coherent Interface", IEEE Standard P1596.3-1996.
- [33] ANSI/TIA/EIA644 standard for Low Level Differential Signalling, 1995.
- [34] "High Performance Parallel Interface, Mechanical Electrical and Signalling Specification", ANSI Standard X3.183-1991.
- [35] A. F. Benner, "Fibre Channel for SANs", McGraw Hill Telecom, 2001.
- [36] H. C. van der Bij et al., «S-LINK, a Data Link Interface Specification for the LHC era », IEEE Trans. Nucl. Sci. vol. 44, no. 3, pp. 398-402, June 1997.
- [37] L. Schmitt et al., "The DAQ of the Compass Experiment", IEEE Trans. Nucl. Sci. vol. 51, no. 3, pp. 439-444, June 2004.
- [38] D. E. McDysan, D. L. Spohn, "ATM, Theory and Applications", McGraw-Hill Series on Computer Communications, 1994.
- [39] P. Steinberg et al. "The Data Collection Modules and ATM-based Event Builder for the Phenix Experiment at RHIC", IEEE Trans. Nucl. Sci. vol. 47, no. 2, pp. 304-308, April 2000.
- [40] K. Anikeev et al., "Event Building and PC Farm based Level-3 Trigger at the CDF Experiment", IEEE Trans. Nucl. Sci. vol. 47, no. 2, pp. 65-69, April 2000.
- [41] J. Kadanbi, I. Crayford, M. Kalkunte, "Gigabit Ethernet", Prentice Hall series in Computer Networking and Distributed Systems, 1998.
- [42] P. Kulnich et al., "The DAQ System with a RaceWay Switch for the Phobos Experiment at RHIC", IEEE Trans. Nucl. Sci. vol. 49, no. 5, pp. 2455-2458, October 2002.
- [43] J. M. Landgraf et al., "The Implementation of the Star Data Acquisition System Using a Myrinet Network", IEEE Trans. Nucl. Sci. vol. 48, no. 3, pp. 601-606, June 2001.
- [44] "LVDS Owner's Manual", 3rd Edition, National Semiconductor, 2004.
- [45] S. M. Rapisarda et al., "Real Time Data Reorganizer for the D0 Central Track Trigger System at Fermilab", IEEE Trans. Nucl. Sci. vol. 50, no. 4, pp. 878-884, August 2003.
- [46] C.A. Siller, M. Shafi, "SONET SDH, A sourcebook of Synchronous Networking", IEEE Press, 1996.
- [47] M. Circella, "Electronics and DAQ of the Antares Neutrino Telescope", Proc. 12th IEEE Real Time Conference, Valencia, Spain, 4-8 June 2001.
- [48] A. Aloisio et al., "An Approach to DWDM for Real-Time Applications", IEEE Trans. Nucl. Sci. vol. 51, no. 3, pp. 526-531, June 2004.
- [49] A. Aloisio et al., "Bus in a New Light", IEEE Trans. Nucl. Sci. vol. 47, no. 2, pp. 309-312, April 2000.
- [50] N. Almeida et al., "The Selective Read-out Processor for the CMS Electromagnetic Calorimeter", Proc. IEEE Nucl. Sci. Symp. Rome, Italy, Oct 2004.