

A High-End Reconfigurable Computation Platform for Nuclear and Particle Physics Experiments

A high-performance computation platform based on field-programmable gate arrays targets nuclear and particle physics experiment applications. The system can be constructed or scaled into a supercomputer-equivalent size for detector data processing by inserting compute nodes into advanced telecommunications computing architecture (ATCA) crates. Among the case study results are that one ATCA crate can provide a computation capability equivalent to hundreds of commodity PCs for Hades online particle track reconstruction and Cherenkov ring recognition.

Nuclear and particle physics studies the elementary constituents of matter and interactions among them. This field is also called *high-energy physics* because many elementary particles don't occur under normal circumstances in nature, but can be created and detected during energetic collisions of other particles in particle colliders. Modern nuclear and particle physics experiments—such as the high-acceptance dielectron spectrometer detector system (Hades, www-hades.gsi.de) and antiproton annihilations at Darmstadt (Panda, www-panda.gsi.de) at the GSI Helmholtz Centre for Heavy Ion Research,

Germany; the Beijing Spectrometer III (BESIII, <http://bes.ihep.ac.cn/bes3/index.html>) at the Institute of High Energy Physics in Beijing; and, at CERN's Large Hadron Collider (LHC, <http://lhc.web.cern.ch/lhc>), the compact muon solenoid (CMS), LHC beauty experiment (LHCb), a Toroidal LHC apparatus (Atlas), and a large ion collider experiment (Alic) achieve their goals by studying the produced particles' emission direction, energy, and mass when the beam hits the target. In such experimental facilities, researchers adopt different kinds of detectors to generate raw data, which are used to calculate and analyze the emitted particles' characteristics after the collision. As an example, Figure 1 shows the exploded view of the Hades detector system.

In high-energy physics, one "event" corresponds to a single interaction of a beam particle with a target particle. An event consists of sub-events that typically represent the information from individual detector subsystems, such as Ring Image Cherenkov (RICH), Mini Drift Chamber (MDC), Time-of-Flight (TOF), and so on (see Figure 1). A detector system often has more than 10^5 signal channels. The delivered data rate, which is the product of the event size and the

1521-9615/11/\$26.00 © 2011 IEEE
COPUBLISHED BY THE IEEE CS AND THE AIP

MING LIU, WOLFGANG KÜHN, SÖREN LANGE, SHUO YANG
AND JOHANNES ROSKOS

Justus Liebig University, Giessen, Germany

ZHONGHAI LU AND AXEL JANTSCH
Royal Institute of Technology, Stockholm

QIANG WANG, HAO XU, DAPENG JIN, AND ZHEN'AN LIU
Institute of High Energy Physics, Beijing

reaction rate, can be a scary number compared to other applications. In Panda, for example, the reaction rate is 10–20 MHz and the data rate is more than 200 Gbytes per second.

Figure 2 shows various experiments by event sizes and reaction rates. Data rates range from 10^7 to 10^{11} bytes/s. This amount of data is too large for the disk or tape storage to record throughout the experiment, which typically last for months. Furthermore, events that are actually of interest to the physicists are rare, and might occur only once within a million interactions. Therefore, it's essential to realize an efficient online data acquisition (DAQ) and trigger system that processes sub-events coming from detectors and reduces the data rate by several orders of magnitude by rejecting the background.

In contemporary facilities, pattern-recognition algorithms^{1–3} such as Cherenkov ring recognition, particle track reconstruction, and TOF processing are implemented as sophisticated criteria according to detector categories. Only those sub-events that possess expected patterns generated by certain particle types and can be successfully correlated among different detectors, receive a positive decision, and are encapsulated in a predefined event structure for mass storage and further offline analysis. Others will be discarded on the fly.

Reconfigurable Computing

The reconfigurable computing paradigm combines software's flexibility and hardware's high-performance using programmable computing fabrics such as field-programmable gate arrays (FPGAs). The fundamental difference compared to ordinary microprocessor computing is that reconfigurable computing lets developers make substantial changes to the data path in addition to the control flow. Although it typically runs at a much lower clock frequency, FPGA-based reconfigurable computing is believed to have a 10–100 times accelerated performance but far lower power consumption compared to general purpose microprocessors (GP-CPU). In physics experiments, FPGA-based solutions have important advantages for implementing pattern-recognition algorithms. First, they have comparatively simple control flows during data processing. Second, the application-specific data path design can result in high performance with the on-chip memory concurrency and fine-grained computation parallelism or pipeline support. In addition, the reprogrammability lets developers change the design according to different experimental requirements.

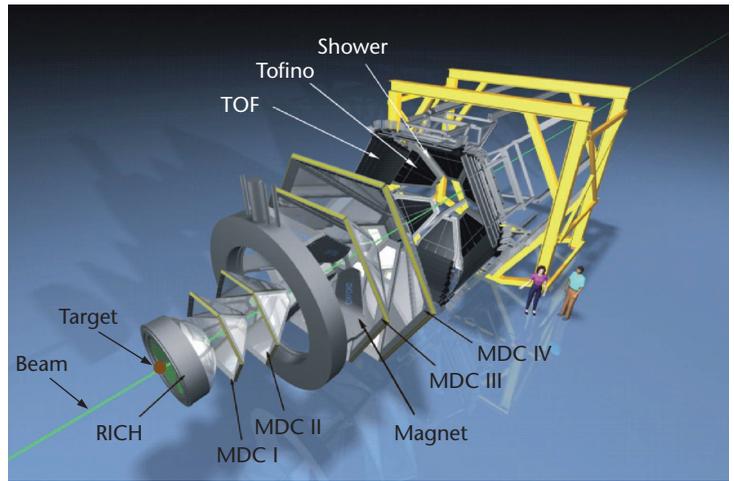


Figure 1. The high-acceptance di-electron spectrometer (Hades) detector system. Researchers use detectors to generate raw data to calculate and analyze particle characteristics after energetic collisions.

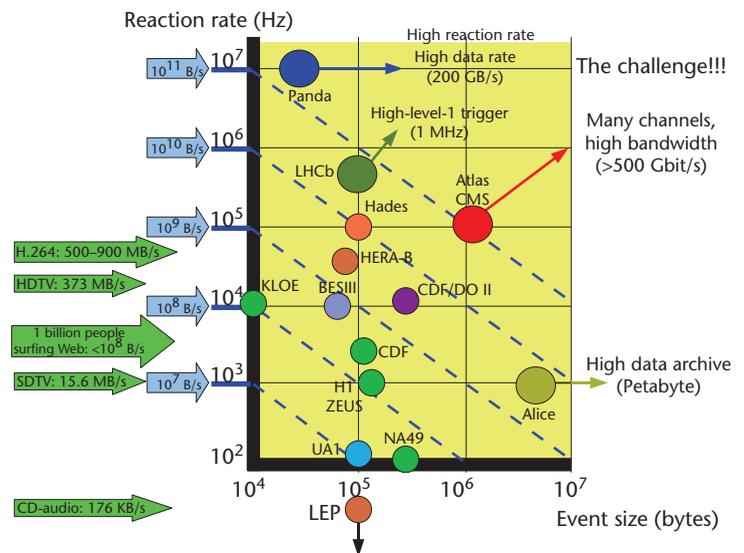


Figure 2. Experiments with different event sizes and reaction rates. Among the projects shown are the high-acceptance di-electron spectrometer detector system (Hades), the Beijing Spectrometer III (BESIII), the Large Hadron Collider beauty experiment (LHCb), compact muon solenoid (CMS), a toroidal LHC apparatus (Atlas), and a large ion collider experiment (Alice).

Motivated by multiple ongoing projects—including the Hades upgrade and the Panda construction—we designed a high-end reconfigurable and scalable computation platform as a general solution. The system is built entirely with commercial off-the-shelf components. We adopted cutting-edge FPGA technologies, as well as high-speed communications, to guarantee high processing capability and channel bandwidth. Easy scalability is an important feature of the platform. To unify application development on

RELATED WORK IN DAQ AND TRIGGER SYSTEMS

Traditionally, developers used modular approaches with commercial bus systems such as VMEbus, FASTbus, and Camac to construct data acquisition (DAQ) and trigger systems for high-energy physics experiments.^{1–4} The bus-based systems containing programmable devices such as field-programmable gate arrays (FPGAs) can interface with PC clusters for hardware/software hybrid processing. However, due to the dramatically increased data rate generated by modern experiments' detector systems, such obsolete technologies no longer meet current requirements. The time-multiplexing nature of the system bus not only deteriorates the data exchanging efficiency among algorithms residing on different pluggable modules, but also restricts the flexibility to partition complex algorithms. Today's networking and switching technologies make it possible to efficiently construct large-scale systems for parallel and pipelined processing. In addition, continuous FPGA development makes it practical to release into the FPGA complex algorithms that were conventionally implemented as software on workstations or embedded processors/DSPs, taking advantage of high-performance hardware processing.

Many commercial and academic projects are now exploring FPGAs' raw computation power for algorithm acceleration. However, they can't be straightforwardly used in physics experiments because of the lack of noise-immune optical links,^{5,6} large communication bandwidth

and memory capacity,^{7,8} or efficient interboard connectivity. Our work takes advantage of point-to-point (P2P) interconnections and modern FPGA technologies to create a hierarchical and scalable computation platform that can optimally interface with other experimental facilities for data acquisition and triggering.

References

1. R. Merl et al., "High Speed EPICS Data Acquisition and Processing on One VME Board," *Proc. Particle Accelerator Conf.*, vol. 4, IEEE Press, 2003, pp. 2518–2520.
2. Y. Tsujita, J.S. Lange, and C. Fukunaga, "Construction of a Compact DAQ-System Using DSP-Based VME Modules," *Proc. IEEE Nuclear and Plasma Societies Real-Time Conf.*, IEEE Press, 1999, pp. 95–98.
3. M. Drochner et al., "The Second Generation of DAQ-Systems at COSY," *IEEE Trans. Nuclear Science*, vol. 45, no. 4, Part 1, 1998, pp. 1882–1888.
4. Y. Nagasaka, I. Arai, and K. Yagi, "Data Acquisition and Event Filtering by Using Transputers," *Proc. Nuclear Science Symp. and Medical Imaging Conf.*, IEEE Press, 1991, pp. 841–844.
5. C. Chang, J. Wawrzynek, and R. Brodersen, "BEE2: A High-End Reconfigurable Computing System," *IEEE Design & Test of Computers*, vol. 22, no. 2, 2005, pp. 114–125.
6. J.D. Davis, C.P. Thacker, and C. Chang, *BEE3: Revitalizing Computer Architecture Research*, tech. report MSR-TR-2009-45, Microsoft Research, 2009.
7. T. Gueneysu et al., "Cryptanalysis with Copacobana," *Trans. Computers*, vol. 57, no. 11, 2008, pp. 1498–1513.
8. O. Mencer et al., "CUBE: A 512-FPGA Cluster," *Proc. Southern Programmable Logic Conf.*, IEEE Press, 2009, pp. 51–57.

the platform, we use a hardware/software code-sign approach to partition functional tasks among embedded microprocessors and modular FPGA cores. Hence, the system design can be largely reused for various experiments with little performance penalty or modification effort. With the uniform hardware architecture and design flow for different projects, manufacturing costs will be largely reduced by mass production, and human resources can be saved to start up development on the platform.

Computation Platform Architecture

To manage the large data rate from detectors, we built a hierarchical network architecture that consists of interconnected compute nodes (CNs). We classify the connectivity as external or internal. The external channels communicate with detectors and the PC farm to receive detector data for processing and forward results for storage and offline analysis. Specifically, they provide optical and Ethernet links. The internal connections

bridge all algorithms or algorithm steps for parallel/pipelined processing. Both onboard I/Os and the inter-board backplane interface function as internal links. We now present a detailed look at the system architecture, starting with the interconnected network and working our way down to the node design.

Network Topology

The Advanced Telecommunications Computing Architecture (ATCA)⁴ standard was established to provide the bandwidth needed for next-generation computation platforms. As Figure 3 shows, a full-mesh shelf backplane can support 2.1 Tbps of data transport when using 3.125 GHz signaling and 8B/10B⁵ encoding. In physics experiment applications, pattern-recognition algorithms are partitioned and distributed in many compute nodes for high processing throughput. Up to 14 nodes can be fitted in one ATCA shelf, and they are mutually interconnected through the backplane. Direct P2P connections provide flexibility

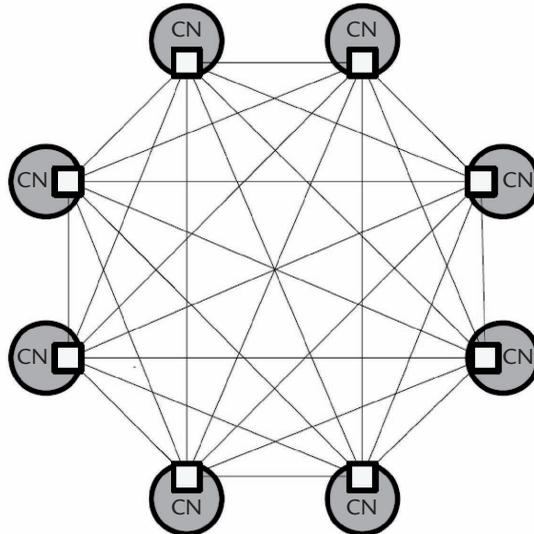


Figure 3. Advanced Telecommunications Computing Architecture (ATCA) crate and full-mesh backplane (only eight nodes shown).

for various network configurations, such as vertical pipelined processing, horizontal parallel processing, or hybrid solutions with more complicated interconnections. This feature enhances the platform's generality for different applications with different network architectures. It also gives developers significant freedom and convenience in partitioning processing logics across multiple boards.

Figure 4 shows the network topology in experimental facilities, where multiple ATCA crates are used to meet high communication and computation requirements. Through bonded optical channels and switches, raw data are dumped from the front-end circuits, which take care of sampling and digitizing analog signals generated by detectors. After that, all data will be processed in the network for pattern recognition, correlation, event building, and filtering. The processing modules are partitioned and reside in FPGA cells. All of these steps constitute the complete computation by communicating through the hierarchical interconnections, including onboard I/Os, inter-board shelf backplanes, and perhaps also the intercrate optical link or Ethernet switching if necessary. Onboard channels provide large bandwidth, while the intercrate switching has more communication overheads and will introduce latency penalty. Thus, trying to group the computation steps with high mutual communication requirements on the same board or next in the same crate is a basic rule for practically implementing the algorithms. After pattern recognition and event selection in the network, most event data are discarded on the fly; only a small part will be labeled as interesting

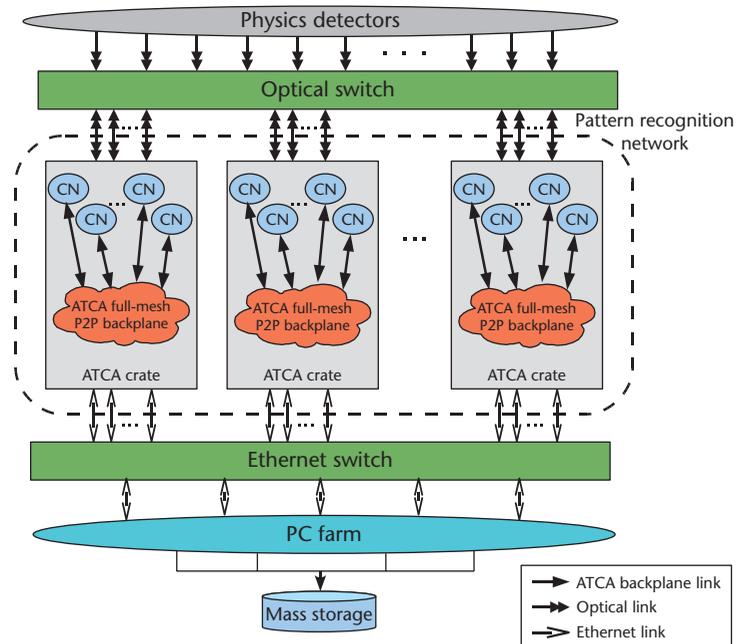


Figure 4. Online pattern-recognition network. Multiple Advanced Telecommunications Computing Architecture (ATCA) crates are used to meet high communication and computation requirements. The processing modules are partitioned and reside in field-programmable gate array (FPGA) cells on the compute nodes (CNs).

and forwarded to the PC farm via Ethernet for storage and in-depth offline analysis.

Compute Node

Our system uses Xilinx platform FPGAs (with hardcore embedded processors) as primary processing components. For the first prototype, we used Virtex-4 FX60. On future products, we might adopt up-to-date generation FPGAs instead.

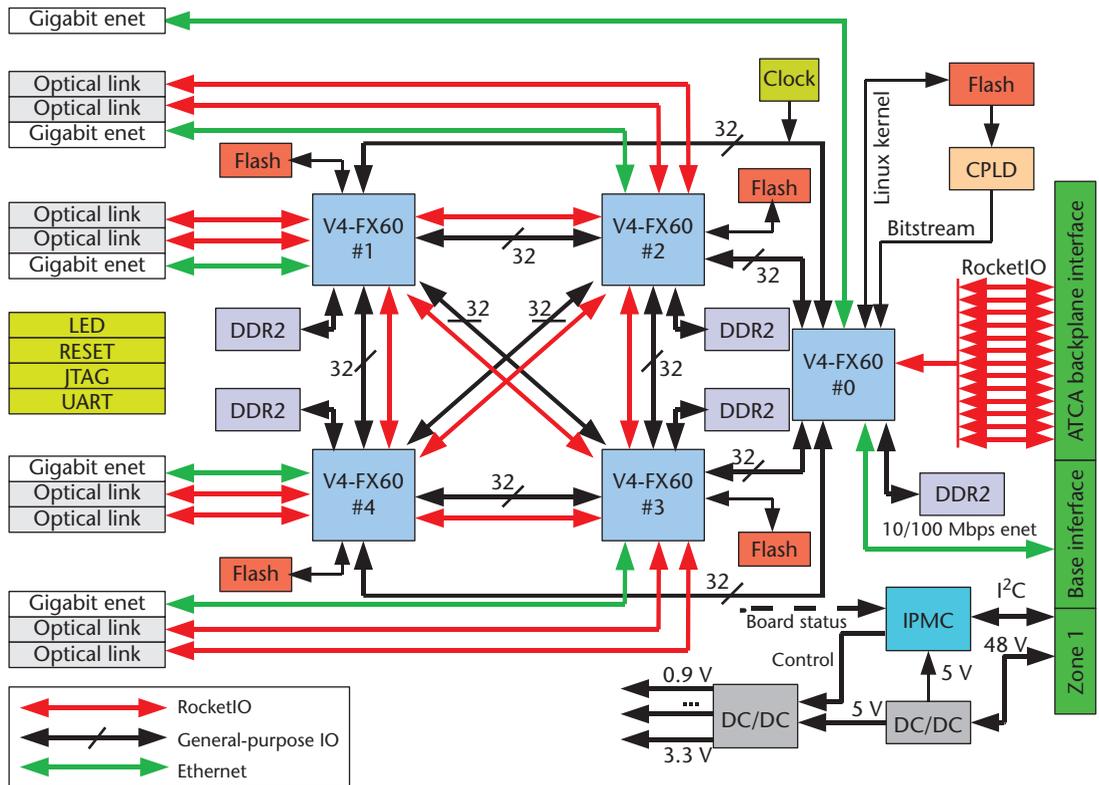


Figure 5. Compute node (CN) schematic. Each board consists of five field-programmable gate arrays (FPGAs): numbers 1–4 are algorithm processors; the fifth, number 0, is a switch interfacing to other CNs through the Advanced Telecommunications Computing Architecture (ATCA) backplane. IPMC stands for intelligent platform management controller; CPLD stands for complex programmable logic device; JTAG stands for joint test action group; and UART stands for universal asynchronous receiver/transmitter.

Figure 5 shows the CN board’s schematic. Each board consists of five FPGAs: numbers 1–4 are algorithm processors; the fifth, number 0, is a switch interfacing to other CNs through the ATCA backplane. Each processor FPGA has two optical links based on RocketIO multigigabit transceiver (MGT). The optical links can run at a maximum baud rate of 6.5 Gbps per channel. In addition, all FPGAs are equipped with one Gigabit Ethernet and 2 Gbytes DDR2 memory each. The total 10 Gbytes memory capacity is mainly for data buffering and large look-up table (LUT) storage.

The processing algorithms and partitions for many future physics experiments are still unclear and might feature different traffic patterns in the network. Thus, to make the board design capable of easily porting high-performance algorithms, all four processor FPGAs are interconnected in a full-mesh topology. The connectivity includes both 32-bit general purpose IO (GPIO) buses and one full-duplex RocketIO link per connection. These processor FPGAs also connect to the switch FPGA with dedicated 32-bit GPIOs. Either

circuit- or packet-switching can be configured to communicate with other CNs in the crate. The 16 RocketIO channels to the backplane feature the bandwidth of 104 Gbps at 6.5 GHz signaling. Besides the switch structure, sub-event data from all four processor FPGAs can be collected in the switch FPGA and conduct event building and filtering. With the onboard P2P interconnections, it’s convenient to partition unexpected algorithms for different experiments and aggregate all five FPGAs as a virtual one with five times the capacity.

Not OR (NOR) flash memories are mounted on the board for operating system kernel and FPGA bitstream storage. A customized intelligent platform management controller (IPMC) add-on card fulfills the ATCA requirements on power negotiation, voltage monitoring, temperature sensing, FPGA configuration check, and so on. Figure 6 shows our first prototype CN printed circuit board (PCB). To meet the dimension requirement, we placed all of the main components on the top, except for five ultra-low-profile small outline, dual inline memory

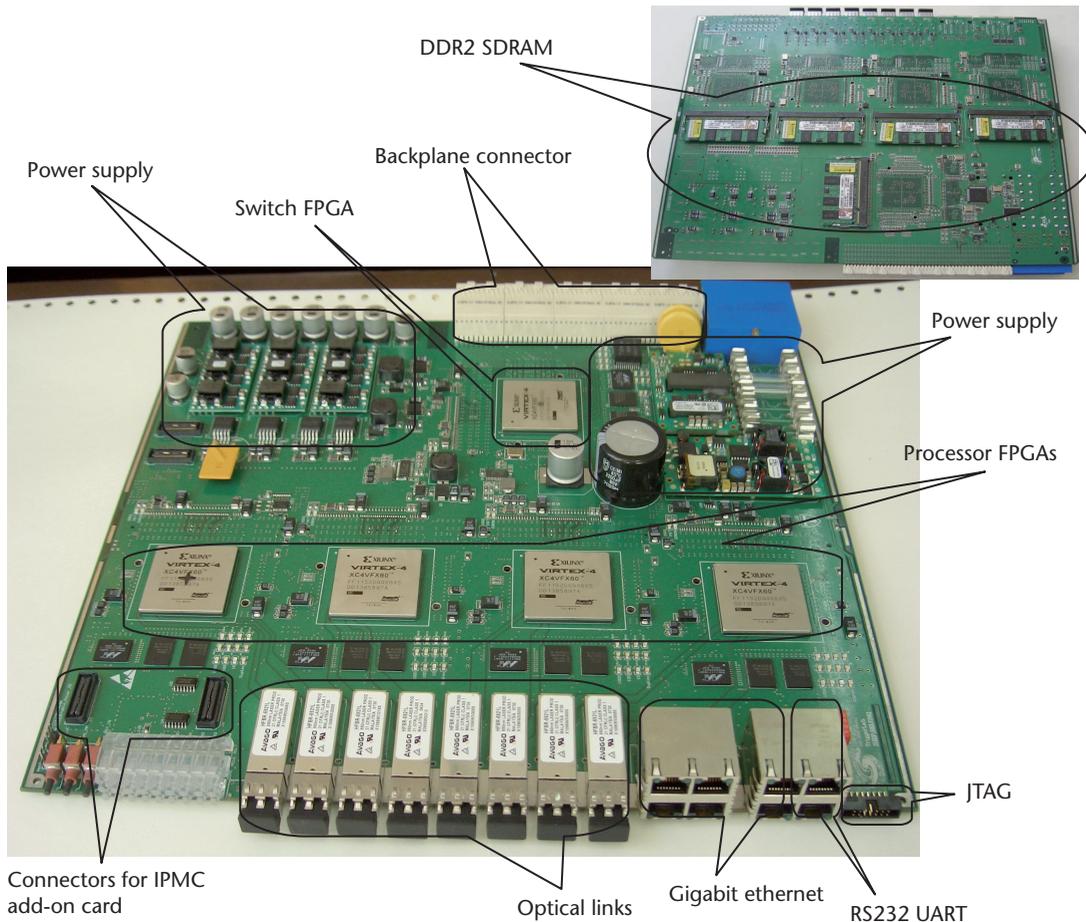


Figure 6. Prototype printed circuit board (PCB) of the compute node (CN). IPMC stands for intelligent platform management controller; JTAG stands for joint test action group; UART stands for universal asynchronous receiver/transmitter; and SDRAM stands for synchronous dynamic RAM.

modules (SO-DIMM) DDR2 synchronous dynamic RAM (SDRAM).

Each CN resides in one of the 14 slots in the ATCA crate. The power budget for each slot is 200 watts at maximum, larger than the worst-case estimation of 170 watts for the CN board. When all 14 nodes are plugged in, such a crate can host up to 1,890 Gbps inter-FPGA onboard connections (GPIOs at 300 Mbytes/s), 1,456 Gbps interboard backplane connections, 728 Gbps full-duplex optical bandwidth, 70 Gbps Ethernet bandwidth, 140 GBytes DDR2 SDRAM, and all computing resources of 70 Virtex-4 FX60 FPGAs.

Hardware/Software Codesign

Based on the computation platform and CNs, we employ a hardware/software codesign approach to ease and accelerate the development for various experiment applications. We now describe the key design issues.

Partitioning Strategy

The DAQ and trigger systems in physics experiments request more features for convenient experiment operations than fundamental data processing. For example, due to temporal and spatial limitations, operators would like to remotely and dynamically reconfigure and control the platform when they're away from the experimental facilities. A friendly user interface also helps physicists easily adjust experimental parameters and monitor the system status. In our platform, designers can utilize hardcore PowerPC microprocessors on the FPGAs to implement versatile control tasks, while locating the performance-critical computation in the FPGA fabric in the hardware.

We have three concrete criteria to partition system tasks:

- All pattern-recognition algorithms should be customized in the FPGA fabric as hardware

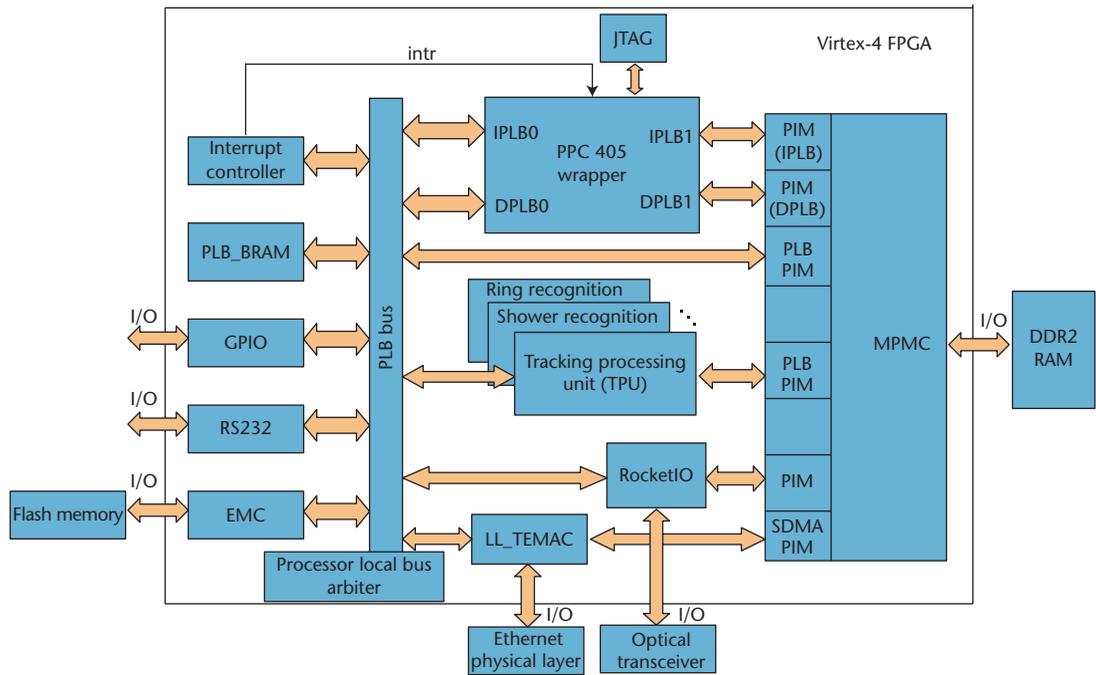


Figure 7. The hardware design is based on the multiport memory controller (MPMC). Compared to standard bus-based designs, MPMC offers direct ports to memory-hungry modules and significantly speeds up memory access. Algorithm coprocessors are customized and integrated in the system design.

- coprocessors working in parallel or pipeline to identify interesting events.
- Slow control tasks should be implemented in software by high-level application programs that execute on top of embedded microprocessors and operating systems.
 - The integrated soft TCP/IP stack in the operating system should be used for Ethernet transmission.

With a reasonable task partitioning strategy, the system can achieve both an optimal data processing performance and design flexibility.

Hardware Design

We adopt the modular design approach to develop the hardware system on the FPGA using hardcore and softcore components. As Figure 7 shows, the PowerPC 405 processor, the multiport memory controller (MPMC),⁶ and other peripherals constitute a complete embedded computer system. Compared to the canonical bus-based design, MPMC provides direct ports to memory-hungry modules and significantly speeds up memory access: incoming detector sub-events are buffered in DDR2 via RocketIO-based optical wrappers and then processed by detector-specific algorithm coprocessors. In this system, the processor local bus (PLB) is used only for low data rate peripheral communications and controls.

For various applications, the system architecture is to be fixed and replace only algorithm engines. It enables design reusability and largely shrinks development time.

Software Design

We ported the open-source embedded Linux kernel (version 2.6) to the PowerPC processor. The soft Linux TCP/IP stack (including UDP/IP) drives the Ethernet transceiving, with commodity PC clusters. Device drivers for standard peripherals can be enabled when Linux is configured, including tri-mode Ethernet, RS232 Universal Asynchronous Receiver/Transmitter (UART), flash memories, and so on. Other drivers for algorithm modules must be customized. Based on the operating system and software development kits, flexible applications can be exploited, ranging from C/C++ and Java programs to high-level scripts. Programs running on PowerPC processors mainly offer user-friendly interfaces for system monitoring and parameter adjustment, drive TCP/IP communications with PC farms, and assist hardware for coprocessing.

One of the software design's main features is its zero budget. Components—including the operating system, the file system generator, the cross-compilation tools, and some benchmark programs—all come from the open source community.

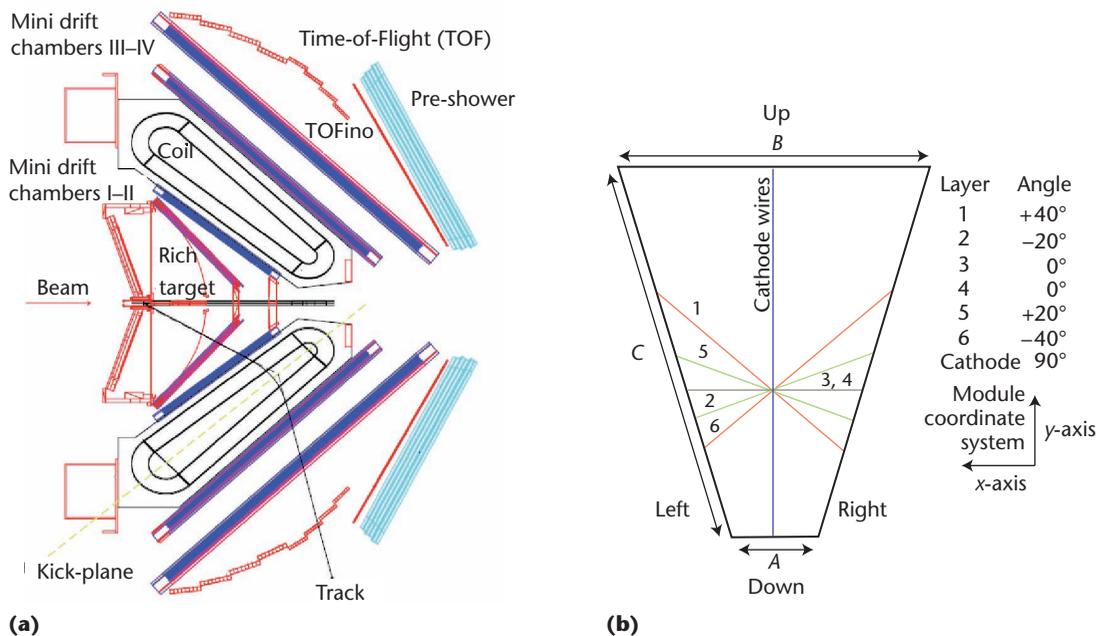


Figure 8. Mini drift chambers (MDCs) in the Hades detector system. (a) A side view of the Hades detector system. (b) One sector of the MDC with six orientation wires.

Case Studies of Trigger Algorithms

Hades is our first experiment in using the computation platform to upgrade the existing DAQ and trigger system for heavier ion reactions and higher processing requirements. The promoted particle reaction rate might reach about 100 KHz, implying a raw data rate up to 10 Gbytes/s (see Figure 2).

Along with the detector construction, we've been developing and evaluating pattern-recognition algorithms on our reconfigurable platform, including Cherenkov ring recognition (for the RICH detector), MDC particle track reconstruction (for MDC detectors), TOF processing (for the TOF detector), and shower recognition (for the electro-magnetic shower detector). All algorithm processors receive readout data and search for certain patterns. Their processing results will be further correlated as well. Only interesting events are assembled and forwarded to the mass storage.

Here, we offer two case studies of algorithm implementations on FPGAs: MDC track reconstruction and the RICH ring recognition.

MDC Track Reconstruction Algorithm

In particle physics experiments, the momenta of charged particles are studied by observing their deflection in the magnetic field. MDC detectors are used to reconstruct the particle tracks entering and leaving the magnetic field to further derive the deflection angle inside it.

The Hades tracking system consists of four MDC modules with six identical trapezoidal sectors. Two MDC layers are located before and two behind the toroidal magnetic field, which is produced by six superconducting coils (see Figure 8a). In first approximation, the magnetic field doesn't penetrate into the MDCs. Thus, particle tracks bend only in the magnetic field; the segments before or behind the coil are approximately straight lines. The two segments can be reconstructed separately with the inner (I-II) and the outer (III-IV) MDC information. Because the basic principle is similar, we focus here only on the inner part.

In the two inner MDC modules, a total of 12,660 sense wires (six sectors) are arranged in 12 layers and six orientations: +40, -20, 0, 0, +20, and -40 degrees (Figure 8b shows one sector). When beam particles hit the target, charged particles are emitted from the target position and go forward through different wire layers in straight paths. Along their flying ways, pulse signals are generated on sense wires close to the tracks with high probability (>95 percent). Hence the sense wires are, in a sense, "fired" by flying particles.

As Figure 9 shows, if each wire's sensitive volume is projected from the target boundary onto a plane located between two inner chambers, apparently the particle passed through the projection plane at the point where all projections from different layers' fired wires overlap. To search for such regions, we treat the projection plane as a 2D

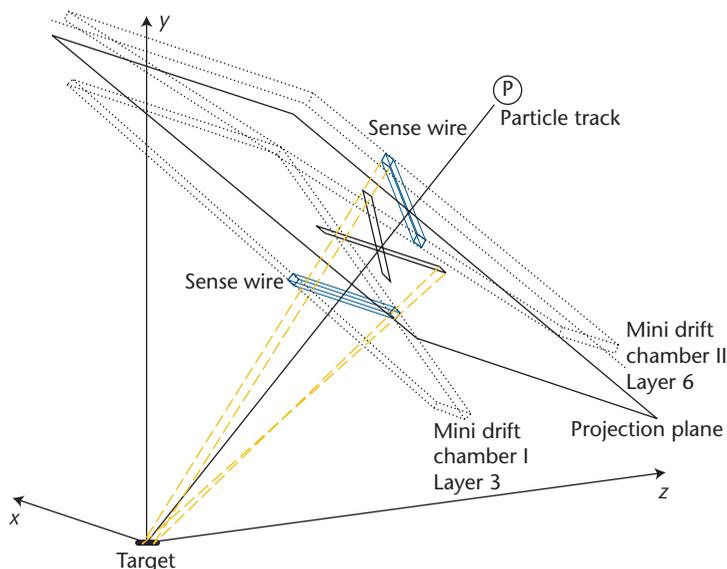


Figure 9. Track recognition and reconstruction in inner MDCs. If each wire's sensitive volume is projected from the target boundary onto a plane between two inner chambers, the particle apparently passed through the projection plane at the point where all projections overlap.

histogram with the projection area as bins (pixels). For each fired sense wire, its projection bins all increase by one. By finding the locally maximum bins whose values are also above a given threshold, track candidates can be recognized and the tracks are reconstructed as straight lines from the point-like target to those peak bins. Figure 10a demonstrates the 2D projection plane for one sector with two passed particles. Figure 10b shows a 3D display of Figure 10a for a single track, where the coordinates of the center peak are recognized as the track's position.

To fit the algorithm on FPGAs, we use a LUT (built offline) to determine which projection plane bins will be touched by the fired wires' projection shadow. We avoid real-time calculation due to the geometrical complexity. For a resolution of 128×256 bins per sector on the projection plane, the projection LUT is about 1.5 Mbytes and is initialized in the DDR2 memory. The tracking processing unit (TPU) is integrated in the system design (see Figure 7). According to the incoming fired wire numbers via optical links, TPU accumulates the histogram of projection on all the bins and searches for peaks where particles most likely passed through. The TPU design features a fine-grained parallel microarchitecture that processes 128 bins (one row) at each time for projection accumulation and peak finding.

Although running at a much slower clock frequency of 100 MHz, a single TPU core can

achieve from 10.8 to 24.3 times speedup in the performance measurements of various wire multiplicities than a single-threaded C prototype program running on an Intel Xeon 2.4 GHz CPU. According to implementation statistics, each TPU core consumes 12.3 percent four-input LUTs, 5.9 percent Flip-Flops and 19.4 percent block RAM (BRAM) resources of one Xilinx Virtex-4 FX60 FPGA. So, this implies that we can integrate at most three TPU cores, in addition to the base system design on each FPGA. Therefore, given the possibility of optimizing the software program and using multicores of modern CPUs, we roughly estimate that one ATCA crate full of CNs can achieve an equivalent processing capability of up to hundreds of commodity PCs.

RICH Ring Recognition Algorithm

The Hades RICH detector identifies dilepton pairs with the Cherenkov light reflected at the mirror. The ring pattern is searched on a detector plane with the resolution of 96×96 pads. The Cherenkov ring features a constant diameter of eight pads on the plane, and the pattern search is executed within a fixed mask region of 13×13 pads (see Figure 11). The hits on a ring with a radius of four pads are added to the value *ring region*. There are two *veto regions* inside and outside the ring region, where the pads are also added. Thus, the ring pattern is identified if both the ring region sum is above and the veto region sums are below their respective thresholds.

Given the constant diameter of ring patterns, the computation challenge is in identifying the position of ring centers. To simplify the algorithm and correlate the RICH pattern with the MDC tracking information, inner particle tracks are introduced in the ring recognition unit (RRU) to point out potential ring centers. A LUT is employed to derive the region of potential ring candidates, converting the coordinate and the granularity from MDC into RICH. With the pointed center, the ring pattern search is efficiently carried out by accumulating the ring region and veto region sums.

The RRU design is currently being implemented. RRUs will also appear as customized coprocessors integrated in the FPGA system design. Taking advantage of the fine-grained parallelism of scanning for rings row by row, each RRU module could achieve a speedup of at least one order of magnitude compared to a software alternative.

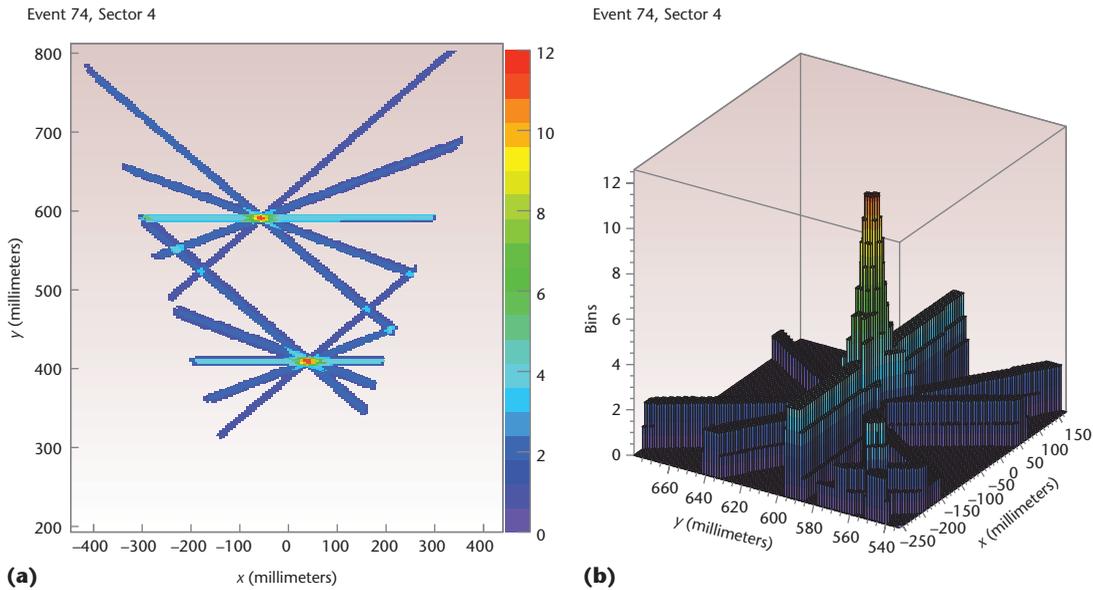


Figure 10. Particle tracks in the projection plane of one sector. (a) A projection plane with two passed tracks. (b) A 3D display of the accumulated bins for a single track.

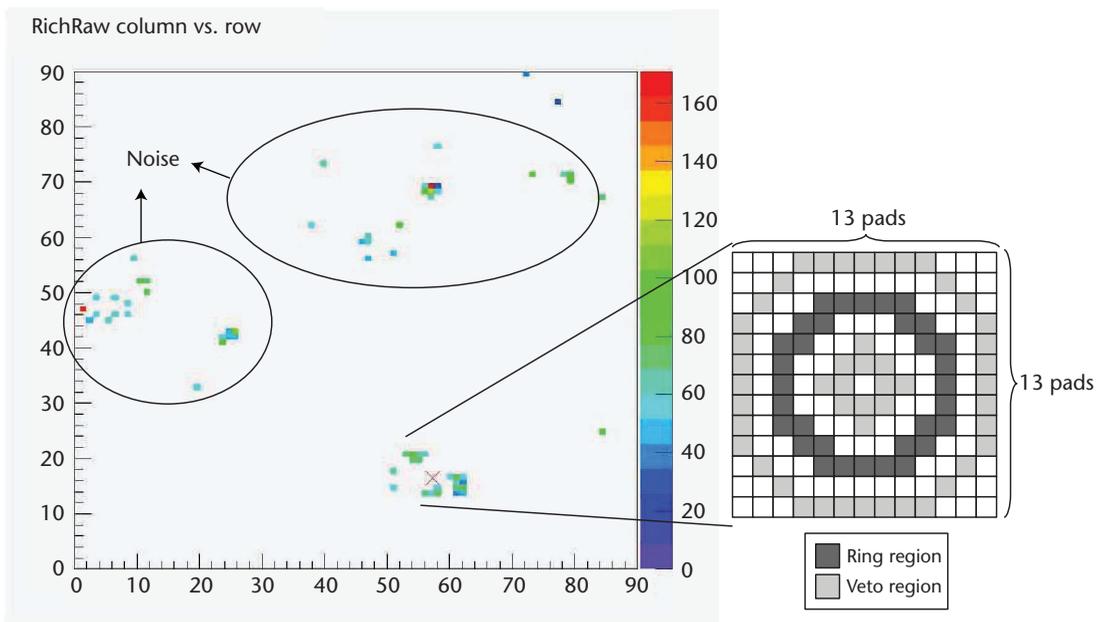


Figure 11. Fixed-diameter ring recognition on the RICH detector. The Cherenkov ring features a constant diameter of eight pads on the plane, and the pattern search is executed within a fixed mask region of 13×13 pads.

Nuclear and particle physics experiments are special applications that are distinguished from common embedded designs such as consumer electronics and mobile devices. The DAQ and trigger systems design addresses great challenges in dealing with the tremendous raw data rate and requiring powerful computation capability. As a general solution for various worldwide

projects, we've presented our hierarchical computation platform based on the ATCA standard and FPGA technologies. The system architecture makes it well suited for hosting trigger algorithms to accomplish particle recognition computation.

As our results from the Hades MDC track reconstruction and RICH ring recognition computation indicate, the platform offers significant data

processing acceleration over canonical software solutions based on commodity PCs.

Given our design's general purpose features, this platform could be promising in other domains—such as weather forecasting, telecom applications, and stock and financial market prediction—to substitute supercomputer software computation with hardware acceleration.

Acknowledgments

Our work was partially supported by the German Federal Ministry of Education and Research (BMBF) under contracts 06GI9107I and 06GI9108I; FZ-Jülich under contract COSY-099 41821475; the Helmholtz International Center for Facility for Antiproton and Ion Research; and Scientific and Technical Cooperation with China (WTZ:CHN) 06/20. We thank Vladimir Pechenov and Björn Spruck for their explanation of the detector-specific algorithms. The authors also thank Xilinx Inc. for the software donation.

References

1. I. Froehlich et al., "Pattern Recognition in the Hades Spectrometer: An Application of FPGA Technology in Nuclear and Particle Physics," *Proc. IEEE Int'l Conf. Field-Programmable Technology*, IEEE Press, 2004, pp. 443–444.
2. M. Traxler, *Real-Time Dilepton Selection for the Hades Spectrometer*, doctoral thesis, Inst. of Physics, Justus Liebig Univ., Giessen, Germany, 2001.
3. C. Hinkelbein et al., "Pattern Recognition Algorithms on FPGAs and CPUs for the Atlas LVL2 Trigger," *IEEE Trans. Nuclear Science*, vol. 48, no. 3, Part 1, 2001, pp. 296–301.
4. PICMG 3.0 Advanced Telecommunications Computing Architecture (ATCA), PCI Industrial Computers Manufacturers Group (PICMG), 2002.
5. A.X. Widmer and P.A. Franaszek, "A DC-Balanced, Partitioned-Block, 8B/10B Transmission Code," *IBM J. Research and Development*, vol. 27, no. 5, 1983, pp. 440–451.
6. *Multi-Port Memory Controller (MPMC)*, Xilinx, 2008; www.xilinx.com/support/documentation/ip_documentation/mpmc.pdf.

Ming Liu is a joint PhD student in the Department of Electronic Systems at the Royal Institute of Technology, Stockholm, and the Justus Liebig University, Giessen, Germany, where he is involved in the construction and development methodology research of a high-end reconfigurable computation platform for nuclear and particle physics applications. His research interests include FPGA-based adaptive computing, high-performance reconfigurable computing, embedded systems, hardware-software codesign, and

network-on-chips. Liu has an MSc in microelectronics from Royal Institute of Technology, Stockholm. Contact him at mingliu@kth.se.

Wolfgang Kühn is a professor at Justus Liebig University, Giessen, Germany. His research interests include experimental particle and hadron physics, as well as applications of programmable electronics in these fields. Kühn has a PhD in physics from Heidelberg University, Germany. He is a member of IEEE and the German Physical Society. Contact him at wolfgang.kuehn@physik.uni-giessen.de.

Sören Lange is a permanent scientific staff member at Justus Liebig University, Giessen, Germany, where he's served in many positions, including deputy technical coordinator for the Hades experiment and deputy computing coordinator for the Panda experiment. His research interests include trigger and data acquisition systems, including specifically applications of transputers, Sharc digital-signal processors, and fast networks (such as Myrinet). Lange has a PhD in nuclear physics from Ruhr University Bochum, Germany. Contact him at soeren.lange@physik.uni-giessen.de.

Shuo Yang is an engineer at Smart Mixed-Signal Connectivity (SMSC) Europe in Karlsruhe, Germany. She previously contributed to the system-on-chip design on FPGAs under a collaboration between Justus Liebig University and the Royal Institute of Technology. Yang has an MSc in microelectronics from Royal Institute of Technology, Stockholm. Contact her at shuo.yang@smc.com.

Johannes Roskoss is an engineer at Esders in Haselünne, Germany, and was a contributor to the physics analysis of the ring recognition algorithm for the Hades experiment's RICH detector. Roskoss has a physics diploma from Justus Liebig University, Giessen, Germany. Contact him at johannes.roskoss@physik.uni-giessen.de.

Zhonghai Lu is a senior researcher in the Department of Electronic Systems at the Royal Institute of Technology, Stockholm. His research interests include network-on-chip/system-on-chip, many-core computing architectures, and cyber-physical systems. Lu has a PhD in electronic and computer systems design from the Royal Institute of Technology, Stockholm. Contact him at zhonghai@kth.se.

Axel Jantsch is a full professor in the Department of Electronic Systems at the Royal Institute of Technology, Stockholm. He leads several research projects, primarily in the areas of system modeling and

network-on-chip. He is the author of Modeling Embedded Systems and SoCs: Concurrency and Time in Models of Computation (Morgan Kaufmann) and is a subject area editor for the Journal of System Architecture. Jantsch has a D.Tech in computer science from the Technical University Vienna. Contact him at axel@kth.se.

Qiang Wang is a joint PhD student at the Chinese Academy of Sciences' Institute of High Energy Physics and the Institute of Physics at Justus Liebig University, Giessen, Germany, where he is involved in designing trigger and data acquisition systems for hadron and particle physics experiments. His research interests include high-speed circuit design, FPGA-based high-performance computing, and sophisticated algorithm design. Wang has a B.Eng in nuclear technology from the University of Science and Technology of China (USTC) in Hefei, China. Contact him at qwang@ihep.ac.cn.

Hao Xu is a nuclear electronics scientist at the Chinese Academy of Sciences' Institute of High Energy Physics. His research interests include low-noise analog circuit, high-speed digital circuit, and FPGA-based system-on-a-programmable-chip (SoPC). He

has a PhD in nuclear technology from the Graduate University of the Chinese Academy of Sciences. He is a member of IEEE. Contact him at xuhao@ihep.ac.cn.

Dapeng Jin is a senior engineer in electronics- and physics-related affairs at the Experimental Physics Center of the Institute of High Energy Physics (IHEP) in Beijing, where his projects include working on the China Spallation Neutron Source's experimental control system. His research interests include beam-related background and implementation of the trigger system for large experiments. Jin has a PhD in nuclear technology from IHEP. Contact him at jindp@ihep.ac.cn.

Zhen'An Liu is a physicist at the Chinese Academy of Sciences' Institute of High Energy Physics (IHEP) and a professor at the Chinese Academy of Sciences' Graduate University. His research interests include experimental physics and its instrumentation, focusing on electronics, trigger, and data acquisition. Liu has a PhD in high-energy physics and nuclear physics from IHEP. He is a member of the IEEE Nuclear and Plasma Sciences Society and the Chinese Electronics Society. Contact him at liuza@ihep.ac.cn.

IEEE computer society

PURPOSE: The IEEE Computer Society is the world's largest association of computing professionals and is the leading provider of technical information in the field.

MEMBERSHIP: Members receive the monthly magazine *Computer*, discounts, and opportunities to serve (all activities are led by volunteer members). Membership is open to all IEEE members, affiliate society members, and others interested in the computer field.

COMPUTER SOCIETY WEBSITE: www.computer.org

Next Board Meeting: 23–27 May 2011, Albuquerque, NM, USA

EXECUTIVE COMMITTEE

President: Sorel Reisman*

President-Elect: John W. Walz;* **Past President:** James D. Isaak;* **VP,**

Standards Activities: Roger U. Fujii;† **Secretary:** Jon Rokne (2nd VP);* **VP,**

Educational Activities: Elizabeth L. Burd;* **VP, Member & Geographic**

Activities: Rangachar Kasturi;† **VP, Publications:** David Alan Grier (1st VP);*

VP, Professional Activities: Paul K. Joannou;* **VP, Technical & Conference**

Activities: Paul R. Croll;† **Treasurer:** James W. Moore, CSDP;* **2011–2012 IEEE**

Division VIII Director: Susan K. (Kathy) Land, CSDP;† **2010–2011 IEEE Division**

V Director: Michael R. Williams;† **2011 IEEE Division Director V Director-**

Elect: James W. Moore, CSDP;* **Computer Editor in Chief:** Ron Vetter†

*voting member of the Board of Governors †nonvoting member of the Board of Governors

BOARD OF GOVERNORS

Term Expiring 2011: Elisa Bertino, Jose Castillo-Velázquez, George V.

Cybenko, Ann DeMarle, David S. Ebert, Hironori Kasahara, Steven L. Tanimoto

Term Expiring 2012: Elizabeth L. Burd, Thomas M. Conte, Frank E.

Ferrante, Jean-Luc Gaudiot, Paul K. Joannou, Luis Kun, James W. Moore

Term Expiring 2013: Pierre Bourque, Dennis J. Frailey, Atsuhiko Goto,

André Ivanov, Dejan S. Milojčić, Jane Chu Prey, Charlene (Chuck) Walrad

EXECUTIVE STAFF

Executive Director: Angela R. Burgess; **Associate Executive Director, Director, Governance:** Anne Marie Kelly; **Director, Finance & Accounting:** John Miller; **Director, Information Technology & Services:** Ray Kahn; **Director, Membership Development:** Violet S. Doan; **Director, Products & Services:** Evan Butterfield; **Director, Sales & Marketing:** Dick Price

COMPUTER SOCIETY OFFICES

Washington, D.C.: 2001 L St., Ste. 700, Washington, D.C. 20036

Phone: +1 202 371 0101 • **Fax:** +1 202 728 9614

Email: hq.ofc@computer.org

Los Alamitos: 10662 Los Vaqueros Circle, Los Alamitos, CA 90720-1314

Phone: +1 714 821 8380 • **Email:** help@computer.org

MEMBERSHIP & PUBLICATION ORDERS

Phone: +1 800 272 6657 • **Fax:** +1 714 821 4641 • **Email:** help@computer.org

Asia/Pacific: Watanabe Building, 1-4-2 Minami-Aoyama, Minato-ku, Tokyo 107-0062, Japan

Phone: +81 3 3408 3118 • **Fax:** +81 3 3408 3553

Email: tokyo.ofc@computer.org

IEEE OFFICERS

President: Moshe Kam; **President-Elect:** Gordon W. Day; **Past President:**

Pedro A. Ray; **Secretary:** Roger D. Pollard; **Treasurer:** Harold L. Flescher;

President, Standards Association Board of Governors: Steven M. Mills; **VP,**

Educational Activities: Tariq S. Durrani; **VP, Membership & Geographic**

Activities: Howard E. Michel; **VP, Publication Services & Products:** David

A. Hodges; **VP, Technical Activities:** Donna L. Hudson; **IEEE Division V**

Director: Michael R. Williams; **IEEE Division VIII Director:** Susan K. (Kathy)

Land, CSDP; **President, IEEE-USA:** Ronald G. Jensen

revised 20 Jan. 2011

