CAPTAN: A Hardware Architecture for Integrated Data Acquisition, Control, and Analysis for Detector Development

Marcos Turqueti, Ryan A. Rivera, Alan Prosser, Jeffry Andresen and John Chramowicz
Electronic Systems Engineering Department, Fermi National Accelerator Laboratory
P.O. Box 500, Batavia, IL 60510-0500, USA

Abstract — **The Electronic Systems Engineering Department of the Computing Division at the Fermi National Accelerator Laboratory has developed a data acquisition system flexible and powerful enough to meet the needs of a variety of high energy physics applications. The system described in this paper is called CAPTAN (Compact And Programmable daTa Acquisition Node) and its architecture and capabilities are presented in detail here.**

The three most important characteristics of this system are flexibility, versatility and scalability. These three main features are supported by key architectural features; a vertical bus that permits the user to stack multiple boards, a gigabit Ethernet link that permits high speed communications to the system and the core group of boards that provide specific capabilities for the system.

In this paper, we describe the system architecture, give an overview of its capabilities and point out possible applications.

Index Terms — **Data Acquisition, Test Stand, Gigabit Ethernet, Semiconductor Detectors, Distributed Computing, CAPTAN.**

I. Introduction

The Electronic Systems Engineering Department in the Computing Division at the Fermi National Accelerator Laboratory has significant experience developing data acquisition systems to support high energy physics research. These systems have been used to characterize individual detector devices and modules for device research. Systems have also been developed to support the commissioning of running experiments such as the Compact Muon Solenoid at the Large Hadron Collider at CERN. In the past, new applications required a significant investment in new hardware, firmware, and software development to support additional requirements. These investments are often costly and frequently are not backward compatible with earlier developments.

CAPTAN is an architecture for distributed data acquisition and processing systems that can be employed in a number of different applications ranging from test stand DAQ (Data Acquisition) systems to high performance parallel computing nodes. In this paper, we describe the CAPTAN system architecture and illustrate system capabilities and limitations. One of the principal goals of this development is to provide a more seamless path to future efforts making efficient use of earlier activities and systems. Such a path is constructed by investing effort into an architecture which is flexible and extendable so that new hardware uses communications, power, and processing resources that are already in place.

Manuscript received Novemberr 14, 2008. This work was supported by Fermi National Accelerator Laboratory operated by Fermi Research Alliance, LLC under Contract No. DE-AC02-07CH11359 with the United States Department of Energy.

M. Turqueti is with Fermi National Accelerator Laboratory, Batavia, IL 60510 USA (phone: +01 630-840-2244; e-mail: turqueti@fnal.gov).

R. A. Rivera is with Fermi National Accelerator Laboratory, Batavia, IL 60510 USA (phone: +01 630-840-3844; e-mail: rrivera@fnal.gov).

A. Prosser is with Fermi National Accelerator Laboratory, Batavia, IL 60510 USA (phone: +01 630-840-3844; e-mail: aprosser@fnal.gov).

J. Andresen is with Fermi National Accelerator Laboratory, Batavia, IL 60510 USA (phone: +01 630-840-3178; e-mail: andresen@fnal.gov).

J. Chramowicz is with Fermi National Accelerator Laboratory, Batavia, IL 60510 USA (phone: +01 630-840-4991; e-mail: chramowicz@fnal.gov).

II. The CAPTAN Architecture

The CAPTAN system is a distributed architecture that is based on core elements known as system nodes. A node is a stack of boards connected together by the vertical bus in which every board in the same node has access to the vertical bus and therefore accessible to each other. There are no limits to the number of nodes that can work together in a system; the only limits are for the number of boards that a stack contains.

The CAPTAN architecture supports two types of data paths, namely, the intra-node and the inter-node data paths. The intra-node communication is achieved by means of the vertical bus that connects all the boards in the same node. The inter-node communications is realized by two different paths, the horizontal bus and the Gigabit Ethernet Link (GEL).

Another key feature of the architecture is the existence of core boards providing the backbone of the node system forming the central part of the hardware. In addition to the core boards (also known as primary boards), there are secondary boards and user boards.

Finally, the software is a big part of the CAPTAN system, and it exists in two levels, the computer to node communications and the network manager. Details about the CAPTAN software architecture can be found in [1].

A. The Vertical Bus

The vertical bus is the main pathway that a board in a node has to communicate directly with another board in the same node. The vertical bus is designed to deal with large amounts of data transported over four identical connectors which are

present on the top and bottom of every CAPTAN core board. A node of two or more boards is assembled by connecting two or more boards together in a stacked arrangement.

Electrically the vertical bus is divided in 12 independent data busses (four with 64 pins, another four with 16 pins and finally, four with 10 pins). In addition to the data bus there is a single 48 pins system control bus and one 16 pins SPI system bus. Power is also distributed to the system over these same connectors providing 3.3V, 5.0V, 12.0V and -12V for the entire node. The power provided by the vertical bus is not necessarily well regulated because the architecture pushes the responsibility for voltage regulation towards the individual boards.

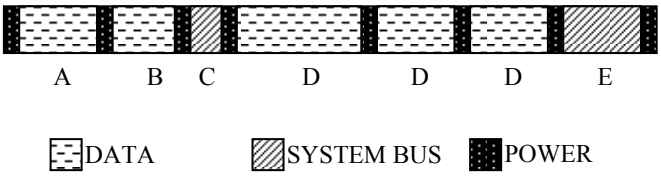

Fig. 1. This figure depicts one of the four identical busses that run vertically through the board where C and E are reserved for the system bus. A, B and D are data buses of widths of 16, 10, 64 pins respectively.

The design of the four buses reflects a high degree of symmetry making the rotational orientation in a stack largely unimportant. Boards can be added to a stack in one of four different orientations. This provides the means to make the best use of board resources in a stack without limitations imposed by the presence of other boards in the stack.

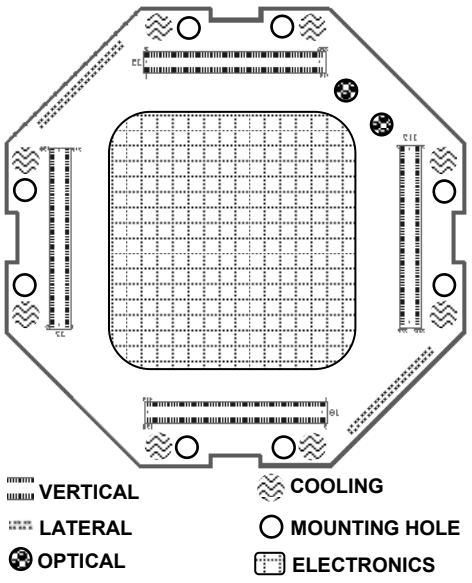

Fig. 2. This figure depicts the baseline design of the any CAPTAN core board with its four vertical bus connectors.

Fig. 1 and 2 provide information about the organization of the bus system. Figure 1 illustrates how the bus is organized on the connector level, while Figure 2 illustrates the distribution of the connectors on the board.

The vertical bus can work in two modes, single-ended mode and differential mode. Since the buses are independent, mixed mode configurations are possible providing flexibility for the operation of devices on the individual boards in a stack. When working in differential mode, the standard supported is LVDS (Low Voltage Differential Signaling) and the number of bits that the specific bus can carry falls by the half of the number of pins available to the bus. Due to the desire to support high speed data communications with robust signal integrity, the CAPTAN system was designed to work primary with differential signals, so its core boards and the bus are designed to support LVDS. The possible vertical bus LVDS configurations for the CAPTAN system are implemented using 32 bits, 8 bits and 5 bits respectively.

The single-ended mode offers a wide range of operation, where each bus can be configured separately to operate as LVCMOS or LVTTL. The use of this mode is desirable when the speed of communications between boards is not an issue. Single-ended mode supports 64 bits, 16 bits and 10 bits data bus definitions.

B. The System Bus

The system bus consists of the system control bus and the system SPI bus. This bus is also distributed in four quadrants across the four inter-board connectors. In addition to the control of the CAPTAN stack, the system bus also supports the programming of devices on the boards of a stack. The controller of the system bus is implemented by a NPCB board (Node Processing and Control Board). Any CAPTAN node is required to contain such a NPCB board. Any NPCB in the node can serve as the node controller but only one node controller can be active at any given time.

The primary task of the system control bus is to carry messages within the node regarding the status of the data bus, ensuring that access to each bus is granted in a safe fashion. Priorities and policies governing bus access are specified on the system bus controller, a firmware block resident in the NPCB.

The system SPI bus is used to distribute the firmware needed to configure programmable logic devices on the boards in the node. The firmware for these devices is delivered from the FPGA on the NPCB. Other signals on the system control bus include a 33MHz reference clock and the node hardware reset signal. The system bus has many pins reserved to accommodate growth of the architecture.

C. The Horizontal Bus

The horizontal bus is a local bus connected only to the electronics of the board itself and does not connect directly to the bus without a bridge. The main reason for the existence of the horizontal bus is to support the connection of secondary boards that can collect data or provide signal conditioning. The horizontal bus can also be used for node-to-node communication, either directly or through a secondary card bridging two nodes.

The configuration of the horizontal bus varies greatly, but the baseline for the digital boards is 32 pins for data communications (differential or single ended), 12 pins for

control signals, clocks and two power supplies. The bus allows also for analog lines for input or output.

D. The Optical Bus

The optical bus is an option to the vertical bus regarding intra-node communications. It is a bidirectional system of high speed lasers for open-air communications capable of providing direct intra-node connections. It is capable of transferring information at up to 1Gbps in a serial fashion regardless of where in the stack the board is located or how many boards are on the stack. There are special requirements that every core board must follow to allow open-air stack communications. The optical transceiver to repeat the laser signal must be placed in specific places on the board or a physical window for the laser to pass.

Fig. 2. , shows where the optical link is located on the core boards.

E. The Node Components and Constraints

One of the key features of the CAPTAN architecture is that it empower the user to develop a wide variety of systems. In the core of this architecture are the system boards that can be divided into core (or primary) and secondary boards. At the same time the architecture is made in such a way to encourage users to design their own boards in cases where the system boards can not fulfill the requirements of the specific application. This is done by providing the user with the foot print of a blank board and the rules that must be followed in order to successfully integrate this new board into the system.

When stacking boards to build a node, there will be constraints, primarily due to resulting limitations on system performance and power delivery, limiting the maximum number of boards that a node can contain.

For example, with every board added to the stack the maximum communication speed that the vertical bus can support decreases due to the effects of cross-talk, signal reflections, capacitive loading, and different delay between lines of the same bus. The maximum number of boards in a node due to these constraints is not easy to predict because it depends on the types of boards being stacked. Experimentally, a limit on the number of boards in a stack is seen to be somewhere between 8 and 12 boards, depending on the types of primary boards which make up the stack.

Power available to operate the stack is limited by the maximum current that the vertical bus can handle and it yields a limit of roughly 100W. This limit is for the whole node, so the number of maximum boards due to power consumption depends on how much power each board consumes, but in practice due to voltage drop in the vertical bus the node should have no more than 10 boards.

III. System Components

The four boards currently considered to be primary boards include the Node Processing and Control Board or NPCB, the Data Conversion Board or DCB, the Power Electronics Board or PEB and the Mass Memory Board or MMB. An important secondary board that is considered a system board is the Gigabit Ethernet Board or GEL. Users are free to design both primary and secondary boards as long as the design rules are obeyed.

A. The Node Processing and Control Board (NPCB)

The NPCB is a primary board that processes information collected by the system or forwards it to another node or computer. The NPCB can handle information from the data buses as well as the system buses. There must always be one NPCB in a node that is responsible for driving the system bus. The NPCB also has two local buses. One is a high speed local bus typically connected to the Gigabit Ethernet Board (GEL). If not used for the GEL, this bus can be used as a general purpose bus. The other secondary bus is a 32 bits single-ended or 16 bits differential general use local bus. The NPCB also provides the gigabit open-air communication system. The system can also be connected to a POF (Plastic Fiber Optic), but if this is done the open-air communication is lost.

The primary component of the current version of the NPCB board is a Virtex-4 Field Programmable Gate Array (FPGA). Seven different versions of this FPGA are supported by the NPCB including the FX12 series devices that possess an embedded power PC core.

When stacking NPCB boards there is a trade off between maximum speed on the vertical bus and maximum number of boards. Table I shows the maximum bus clock speed obtained experimentally with one node (configured with different numbers of boards) for which no errors were observed in intra-node communications on the vertical bus. The bus was tested in a loopback configuration in both single-ended and differential modes. When in differential mode these values were experimentally obtained with double termination with the exception of a two boards stack.

TABLE I
MAXIMUM BUS CLOCK SPEED GUARANTEED

Number of Boards on the Node	Maximum Vertical Bus Clock Speed in MHz (single)	Maximum Vertical Bus Clock Speed in MHz (differential)
2	200	340
3	150	280
4	120	240
5	80	200
6	66	150
7	33	125
8	-	100
9	-	66
10	-	33

The single ended version of the test works with the top NPCB generating a 64 bits word and presenting the data to one of the four available 64 bits data bus. The data is then transmitted on this bus to the last board on the stack. The last board loops back the information through a second 64 bits wide bus and the information is presented back to the top NPCB. At this point the data is compared and if the received data does not match the transmitted data, an error is recorded.

The electrical protocol used for this test was LVCMOS operating at 1.8V.

The differential test works in the same way, except that the buses were 32 bits wide and there were 100 Ohm terminations in the node placed on the receivers at the two extremes of the stack (one at the top and the other at the bottom). The electrical protocol used for this test was LVDS working at 2.5V. In both the single-ended and differential mode tests, if no errors were recorded for a transmitted data payload of 10GB, the test was considered to have passed.

Simulation was used in order to help the analysis of the data integrity on the bus and it matched within 20% of measurements from the bus. The basic parameters measured and compared with simulation included eye pattern opening, delay within a bus and propagation time. While it is not the scope of this work to analyze in depth the results of the electric tests performed on the bus, the results seem to indicate that the limits presented on TABLE I can not be improved much further. The board to board propagation time is a very important parameter however and its measurement provided a mean of 80ps board to board propagation time without adding delays due to the I/O buffers that vary depending of the FPGA grade being used.

Also is important to mention that the normal power consumption of this board is 1.5W when not providing power to any secondary board.

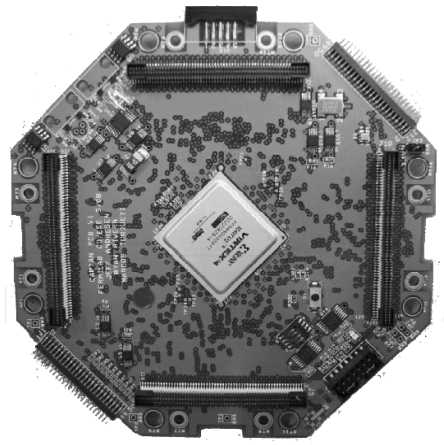

Fig. 3. Shows a picture of the NPCB board.

B. Data Conversion Board (DCB)

The Data Conversion Board is the primary system board used for data acquisition. This board includes Analog to Digital Converters (ADCs) and Digital to Analog Converters (DACs).

There are three different ADCs on the board, providing varying conversion speeds. TABLE II provides some of the details for the different converters. The medium speed ADC hardware includes an operational amplifier in each channel, where gain and offset can be adjusted digitally.

Also available on this board are two DACs, one fast and one slow. A summary of the capabilities of the DACs is presented in

.

TABLE II
SUMMARY OF THE ANALOG TO DIGITAL CONVERTER CAPABILITIES OF THE DCB BOARD

ADC TYPE	# of channels	Resolution	Sampling
ULTRAFAST	2	8 bits	1.6 Gsps
MEDIUM	8	12 bits	65 Msps
SLOW	12	12 bits	0.3 Msps

Access to the ADCs and DACs is provided on three lateral connectors on the board, with exception of the ultra fast ADC that has special dedicated coaxial connectors. The board also provides a bridge to the vertical bus through a fourth lateral connector.

TABLE III
SUMMARY OF THE DIGITAL TO ANALOG CONVERTER CAPABILITIES OF THE DCB BOARD

DAC TYPE	# of channels	Resolution	Sampling
FAST	2	12 bits	500 Msps
SLOW	32	16 bits	0.3 Msps

This bridge is capable of level translation and is meant to provide access to the vertical bus to external peripherals.

In addition to the above features, the DCB board also provides five adjustable regulated voltages through the lateral connectors. TABLE IV provides the capabilities of the voltage regulators available on this board to the lateral connectors.

TABLE IV
SUMMARY OF THE BOARD REGULATION CAPABILITIES

REG TYPE	# of channels	Range	Max current
POSITIVE	3	1.0V to 5.0 V	2000mA
NEGATIVE	1	-2.5 to -5.0V	200 mA
HV	1	50 to 250 V	2 mA

To configure and control devices on the DCB, interfaces with the system bus by means of four CPLDs (Complex Programmable Logic Device) are provided (one for each of the four system bus quadrant connectors), making available to the system bus parameters such as sampling rate, voltage references, channel gain, offset and voltage levels.

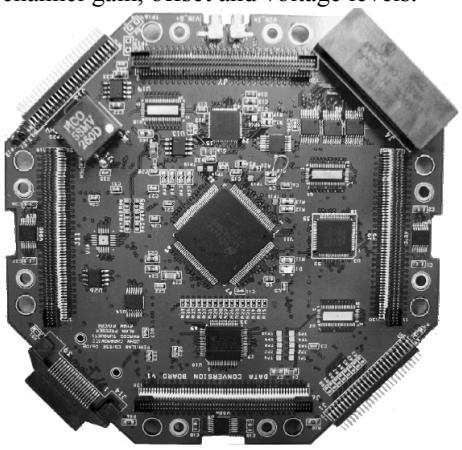

Fig. 4. The picture shows the DCB board, its four lateral connectors are visible in to the angled edges of the board.

Fig. 4. shows the top portion of the DCB board (the 1.6Gsps ADC is at the center of the board).

C. The Power Electronics Board (PEB)

The Power Electronics Board is a primary board that provides regulated power supply for the four power lines available on the vertical bus, although the stack can be powered externally without the power board.

The power board provides regulated monitored power and protection to the bus. The power board also provides 8 IGBT high power drives and 8 MOSFET medium power drivers in order to control devices that require high current switching capability. This board can provide a maximum combined power of 100W to the vertical bus and additional 20 W for the onboard switches.

D. The Mass Memory Board (MMB)

The Mass Memory Board is a primary board that provides up to 1GByte of DDR2 memory. The objective of this board is to provide a high speed local mass memory to the system. The design of the board supports standard SO-DIMM DDR2 memory cards and every card is connected to a carrier CAPTAN board that is connected direct to the 64 bits bus and they can only connected to the bus when this is operating in single ended mode LVTTL 3.3 V. This board is still under development and more results on its performance will be provided in the future.

E. The Gigabit Ethernet Link (GEL)

The GEL is a secondary board providing gigabit communication between nodes or between a node and a computer. This board is the main external interface of the CAPTAN system and can communicate directly with any computer with 1000BASE-X network capabilities. The board is designed to work with Ethernet protocol 10/100/1000 and to use UDP/IP as the communication protocol.

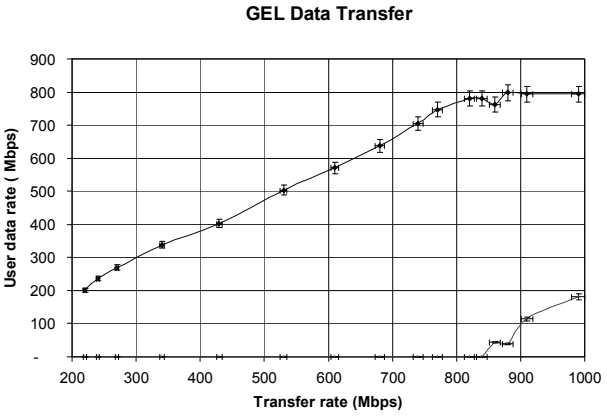

Fig. 5. This graph shows how fast the user can transfer information before it starts to loose data, the top curve refers to data successfully transferred and the bottom curve to data lost.

Although the board is capable of connecting using the IEEE 802.3ab (1000BASE-X) protocol it can not send pure user data at this speed due to the addition of several layers of protocol, maximum packet size limitations and the particular hardware used. The board can however send pure user data up to 800 Mbps as illustrated in Fig. 5. The performance demonstrated in Fig. 5 was obtained using the UDP protocol with the packet size set to 1466 bytes with the PHY running 1000BASE-X.

The power consumption of the GEL board is 1.25W. Any NPCB in a stack can support a GEL board. The GEL board can currently only interface with the node through the NPCB board and the NPCB FPGA firmware must include the gigabit Ethernet controller firmware as the GEL board itself contains only the gigabit PHY.

F. User Boards

Guidelines are provided in order for users to design their own boards compatible with the CAPTAN system. Either primary or secondary boards may be designed but secondary boards tend to have less overhead than primary boards.

In order to design a primary board there are three main aspects that must be taken in consideration. The first one are the vertical bus communication and electrical rules that must be followed. For example, a primary board must have buffers to isolate itself from the vertical bus. Second, power limits must be obeyed (the board can not consume more than 12W of power). Finally the designer must follow the mechanical rules for the primary board type, which will guarantee that cooling channel, optical link and bus interfaces will match the existing system.

For the design of a secondary board, the rules are less restrictive as the only mechanical constraint that the user will face is on the lateral bus connection. Also, if the user is utilizing power from the stack through the lateral bus, it can not exceed 3W.

IV. SYSTEM APPLICATIONS

There are a number of possible applications for the CAPTAN system, including data acquisition systems, data processing systems, and mixed applications. The flexible nature of the CAPTAN architecture makes possible topologies which can be individually suited to the application. Some examples of topologies to support different applications will be illustrated next.

A. Single Core Board Application

The powerful FPGA of the NPCB makes it possible to configure a CAPTAN system with Gigabit Ethernet service using a single primary board and a GEL board. Depending on the FPGA used on the NPCB board, up to 400 I/O pins that can be configured as 400 LVCMOS or LVTTL signals or 200 LVDS signals are available as user I/O. This option provides ample support for many applications that do not require analog capabilities. One such application from detector R&D for high energy physics is the construction of test stands for the evaluation of digital pixel readout ICs. Since the NPCB board can be assembled with an FPGA that contains embedded PowerPC core, applications which benefit from embedded software solutions can be supported. As the basic firmware that is required to be implemented on the NPCB FPGA occupies less than 10% of the FPGA, there are considerable resources available for the implementation of application

specific firmware (such as data compression or digital signal processing blocks).

B. Multiple Board Application

Stacking multiple primary CAPTAN boards creates a more powerful node and greatly expands the capabilities of the system. Such a node may be capable of dealing with both analog and digital information (if it includes a DCB as one or more of the up to 10 primary boards). Depending on the number of NPCBs in the stack, up to 10 GEL boards can provide networking access to multiple boards in the node. Nodes of this type find applications as test stands for pixel detector readout devices with analog readout architectures. These nodes can be used as laboratory bench test systems for characterizing individual modules for a pixel detector.

TABLE V shows three of many possible configurations that a multiple board node can assume and some key capabilities that these nodes would offer. The data rates between the boards of a node and between other networked elements (e.g. a personal computer or another CAPTAN node) are listed for different stack compositions (i.e., different combinations of NPCBs and DCBs in a node).

TABLE V
NODE PERFORMANCE IN DIFFERENT CONFIGURATIONS

Total Number of Boards	Number of Gsps ADC channels	PowerPC Cores	Intra-node data rate transfer	Inter-node data rate transfer
8	0	8	19 Gbps	6.4Gbps
8	14	1	19 Gbps	0.8Gbps
8	8	4	19 Gbps	3.2Gbps

Three different configurations of the node are exposed, the first one counts with 8 NPCBs boards, the second one 7 DCBs boards and one NPCBs and the third one 4 NPCBs boards and 4 DCBs boards.

An extra resource that a stack can utilize is to use its GEL boards for communication within the stack, this can greatly expand data rates within the node.

Fig. 6 shows an actual stack composed of three NPCB boards, on the top of the stack is a JTAG connector used to download the initial firmware to the FPGA. While NPCB boards inserted on the stack must initially be configured with the basic firmware, the firmware can be updated later over the Ethernet connection. In addition, any DCBs in the stack can also have firmware updates applied in-system by using the FPGA on an NPCB in the node as a programming device. The bit stream is transferred over the vertical bus and delivered to a quadrant CPLD.

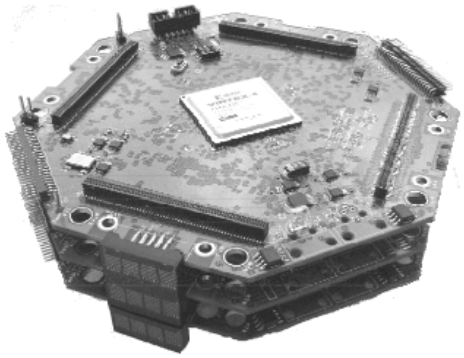

Fig. 6. Node configured with 3 NPCBs boards.

C. Distributed Network Applications

Distributed systems with multiple CAPTAN nodes can be configured using the Ethernet capabilities of the nodes. There is no limit in the number of nodes that can be interconnected making possible for such a configuration to deal with very large amounts of data at very high speeds.

An example of such a networked application is an extension of the pixel detector readout systems used in test benches to a pixel-based tracking readout system for supporting a test beam telescope for particle physics research. Due to the scalable nature of the architecture, this design can be expanded to include thousands of detectors in a straightforward manner to support the data acquisition needs of a larger pixel tracking system.

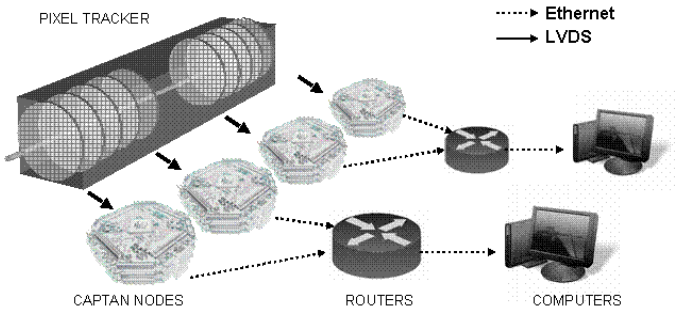

Fig. 7. Applications such as pixel tracker readout for high energy physics are possible.

Fig. 7 shows such an application, utilizing commercially available components such as routers, switches and Ethernet cards compatible with the CAPTAN system nodes. Such a system makes the best use of cost-effective off the shelf components with application specific CAPTAN components. A software architecture suitable for implementing such networked applications is described in [1].

Another possible application for a networked CAPTAN application is as a computing farm. Such an application could support special purpose computational tasks, possibly requiring digital signal processing requirements with more general purpose computing tasks, in a parallel computing implementation. Another example would be to evaluate the temporal evolution of a network of connected elements in a discrete event simulation. The nodes can be configured (in firmware or software) to evaluate the individual states of the

connected elements. The communications needed to update the state of one element to reflect a change in state of another element would be conveyed over the vertical bus in a stack (for local connections) or over the Ethernet links (for more remote connections).

V. Conclusion

The CAPTAN system is powerful and flexible data acquisition system architecture with its roots in test beam and pixel detector tracking applications for high energy physics. However, due to the flexible and expandable characteristics of the system, it has found applications outside the high energy physics world. The CAPTAN architecture is novel in that the systems scale by taking advantage of the node approach for easy growth. It also supports multi-process and distributed computing, and provides added capabilities through convenient interfaces to analog instrumentation.

Examples of CAPTAN systems are being deployed to implement diverse applications that range from detector test beam and test stand systems to multi-core DSP applications. It is hoped that the diverse set of applications supported will give rise to a vibrant community of collaborators leading to the efficient reuse of development products.

References

[1] R. A. Rivera, M. Turqueti, A. Prosser, S. Kwan, "A Software Solution for the Control, Acquisition, and Storage of CAPTAN Network", procedures NSS2008.

[2] IEEE Standard for Compact Embedded-PC Modules P996.1 Draft D1.00, pp. 1-16, 20 Oct. 1993.

[3] IEEE P802.3ab 1000BASE-T

[4] ESE PCI Test Adapter, FERMILAB document CD 1590-V1.