

Heterogeneous Processing Platform FPGA Solutions for the IBM BladeCenter™

White Paper: 200905.01
Prepared by Annapolis Micro Systems, Inc.
May 20, 2009

1. Introduction

Faster processing. Better application performance. Do more work and do harder work. Do the work better and do the work faster. Use less power. Take up less space. Pay less to buy, pay less to program, and pay less to operate. These are the criteria for a successful high performance computing platform.

A successful Heterogeneous Processing Platform provides multiple types of processing elements, high levels of interconnectivity between the processing elements, multiple fast connections between the outside world and the processing system, copious amounts of memory easily accessible by the processing elements, ample power and cooling, and the means to quickly and effectively program the applications.

Recognizing the increased demands for high-performance computing, Annapolis Micro Systems has selected and targeted their WILDSTAR™ 5 Blade architecture for the IBM BladeCenter computing platform. With this approach, system designers can complement traditional microprocessor blades such as the IBM Cell Processor, IBM PowerPC, AMD Opteron, etc., with FPGA computing blades in proportions that make sense given the algorithms required and the processing time specified. This type of “hybrid computing” will ultimately accelerate the processing of complex applications. Multiple WILDSTAR 5 Blades can be installed into one or more IBM BladeCenter chassis to provide the wide scaling necessary to help solve computational intensive problems.

The IBM BladeCenter is a powerful, thoroughly integrated platform that provides the highest thermal and power capacities in the industry today. It's intelligent system design with multiple layers of redundancy, single point of control management, powerful processor blades and the ability to insert additional blades when needed, provides the perfect foundation for a heterogeneous processing platform to run compute intensive applications.

The WILDSTAR 5 for IBM BladeCenter is the newest member of Annapolis Micro Systems' WILDSTAR family of FPGA-based computing products. This new innovative blade approach allows system designers to mix and match traditional microprocessor blades with FPGA and processor computing blades within the same BladeCenter chassis. The WILDSTAR 4/5 family of I/O cards provides a wide variety of world class analog to digital inputs and digital to analog outputs, as well as additional high speed communication access into the Blade Center.

Within the WILDSTAR Blade itself, the architecture provides the ultimate in scalability and modularity to provide solutions for the most complex real-time signal processing or advanced networking problems.

Figure 1 shows the WILDSTAR 5 for FPGA Blade.




Figure 1: WILDSTAR 5 FPGA Blade

2. *FPGA Blade Architecture*

The WILDSTAR 5 FPGA Blade architecture is best illustrated in terms of Computational Processing Elements (CPEs), I/O Processing Elements (IOPEs), multiple memory ports, an extensive system of Serial Gigabit interconnectivity, up to 4 pluggable external I/O modules, and access to up to two slots of the BladeCenter midplane for each WILDSTAR 5 board, as shown in Figure 2.

IOPEs are chosen at manufacturing time, and are not removable. Current IOPE choices cover a wide range of parts from Xilinx's Virtex-5 FXT, LXT and SXT families. Refer to Figure 2. The IOPE supports up to three external DDR2 DRAM memory ports.

The CPEs are on Computational Pluggable Modules (CPMs). These modules are ordered separately, and can be mixed and matched, and removed and replaced by the user. Current modules cover the same range of Xilinx Virtex-5 FXT, LXT and SXT families available as IOPEs. The CPMs have up to seven external memory port options of either DDR2 DRAM, DDR2/QDR2 SRAM, or SIO/CIO RLDRAM. Up to six CPMs can be populated on the FPGA Blade baseboard. Additional types of CPMs will be available in the future.

The FPGA Blade architecture includes three Serial Gigabit I/O (SGIO) switches. First, a high-speed, protocol agnostic Crossbar switch allows for high-bandwidth, full-duplex communication paths between all processing elements, front-panel I/O, and the IBM BladeServer Midplane. Secondly, a standard PCI-Express Gen 2 switch is also on-board for PCIe switched protocol connectivity between all CPEs and three PowerPC embedded processors. Finally, a Gigabit Ethernet Switch allows for gigabit Ethernet switched protocol connections between the PowerPCs and the IBM BladeServer Midplane. All three SGIO switches are software configurable via the Host PowerPC embedded processor running Linux.

In addition to the onboard Host Block PowerPC processor, two additional embedded PowerPCs are also available for user applications running Linux. These embedded processors with their direct connection to their respective IOPEs could be used to process complex protocols such as Infiniband.

External sensor I/O connections are provided by the family of Annapolis WILDSTAR 4/5 Front Panel Mezzanine Cards. A broad range of A/D conversion products ranging from 130 MSps to 5.0 GSps is available. D/A conversion products range from 600 MSps to 4.0 GSps. High-speed serial communication interfaces are also available over CX-4 copper/fiber optic media, Tri-XFP 10 Gigabit fiber optic I/O, and Quad SFP+ 10 Gigabit fiber optic I/O. These SGIO mezzanine cards support standard protocols like Infiniband, 10G Ethernet, and Serial FPDP. Each FPGA Blade baseboard can accommodate up to two mezzanine cards on a single-wide Blade and up to four mezzanine cards on a double-wide Blade.

The WILDSTAR 5 FPGA Blade is fully compliant with the IBM Blade Management Controller (BMC). The enhanced BMC for the FPGA Blade is a flexible service processor that provides support for the following functions:

- Intelligent Platform Management Interface (IPMI)
- Operating System
- Power control and advanced power management
- Reliability, availability, and serviceability (RAS)
- Continuous health monitoring control
- Configurable notification and alerts
- Event logs that are time-stamped and stored in nonvolatile memory
- Remote power control
- Remote firmware update and access to critical FPGA blade settings

Built-in environmental management automatically monitors and controls temperature and power, which allows the user to concentrate on the task at hand. The FPGA Blade permits all available power to be used, but remain managed and controlled.

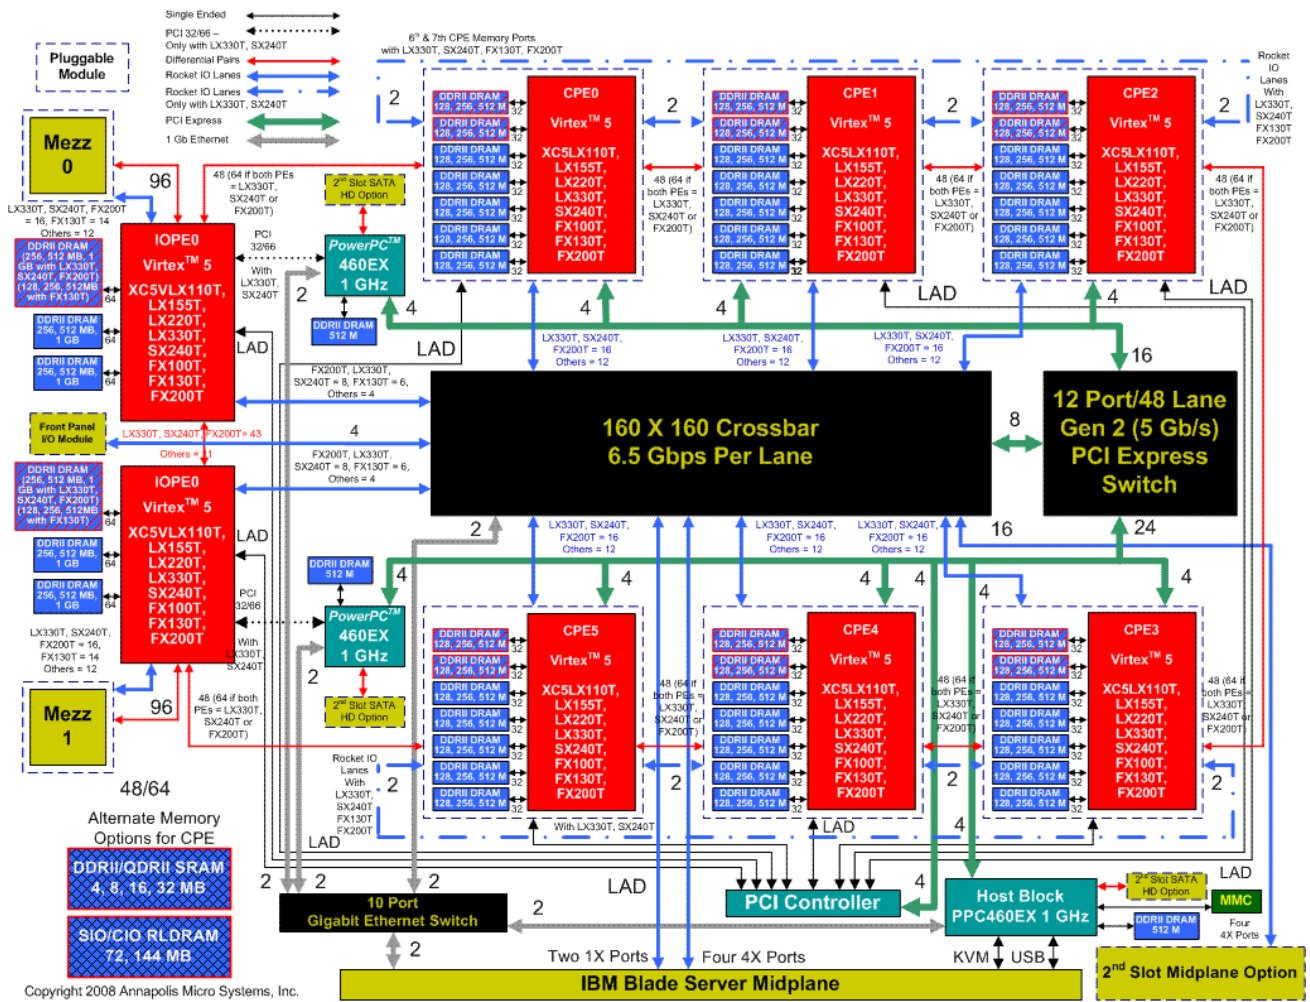


Figure 2: WILDSTAR 5 FPGA Blade Block Diagram

3. Scalable FPGA-based Real-time Signal Processing Platform

3.1 Radar Receiver Application

Figure 3 illustrates a dual-channel FPGA-based Radar Receiver implemented on the Annapolis WILDSTAR 5 FPGA Blade. It highlights the key functional blocks required to implement such a real-time processing system. Two front panel A/D mezzanine cards are used to digitize the incoming RF channels up to a sampling rate of 2 GSps. The next stage includes a Digital Down Converter (DDC) necessary to convert the digitized real signal to a basebanded complex signal centered at zero frequency. Downstream from the DDC, low-pass FIR filters pass only the desired signal and perform anti-aliasing filtering prior to decimation.

In parallel, time-domain overlapping FFTs can be performed along with a threshold detection operation that is used to implement a signal detection function. Fixed-length or variable-length FFT cores can be instantiated to allow frequency resolution to be traded for a faster detection response. For example, when a signal detection FFT is configured for 4K points, the core can perform 50,000 FFTs per second, whereas a 16K point FFT can achieve 12,000 results per second. In addition, fast ultra-long FFTs can be implemented using the Cooley-Tukey algorithm, reducing the algorithm to two smaller FFT cores ($N_1 \times N_2$), Corner-Turn functions utilizing Matrix Transpose IP cores and external DDR2 SRAM banks all resident on a single Virtex™-5 CPE module (CPM).

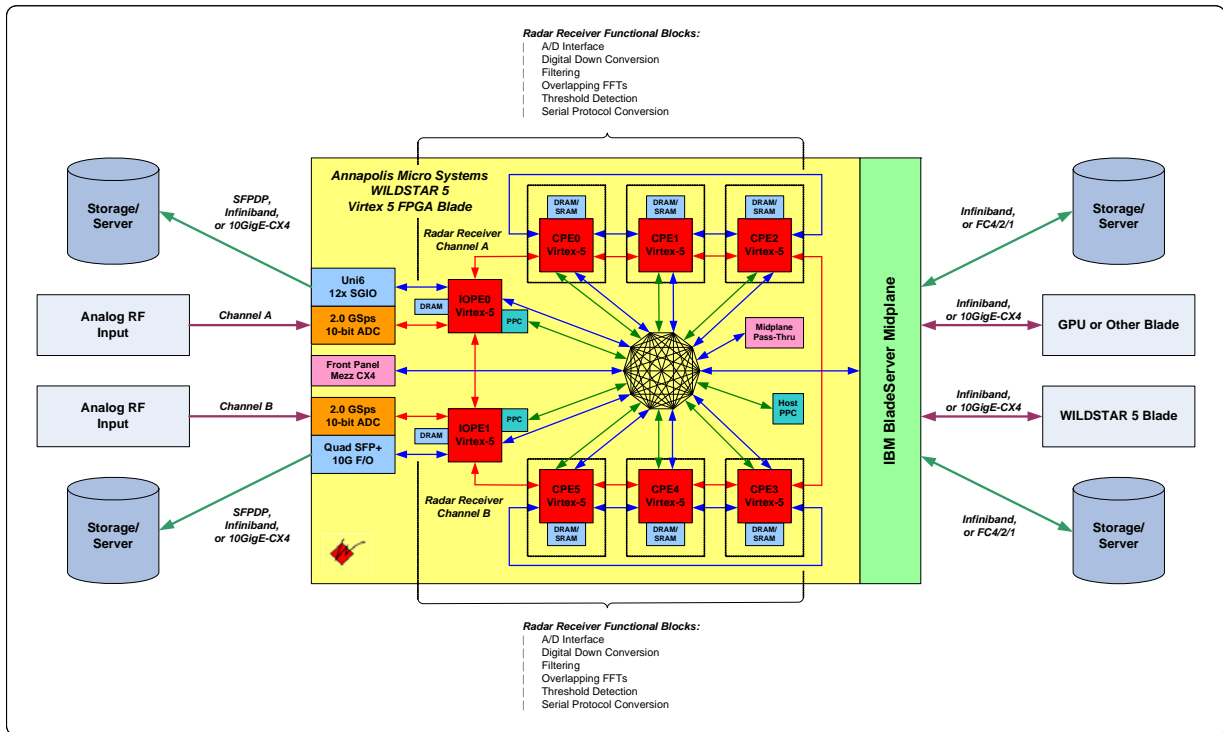


Figure 3: FPGA Blade Radar Receiver Application

The spectrum information, data stream, time-stamping, etc. are aggregated and sent over a single network interface to a storage or server system. IP cores are provided to implement industry-standard protocols such as Serial FPDP, Infiniband or 10 Gigabit Ethernet over copper or fiber media. Network connectivity can also be managed through the BladeServer midplane and forwarded to other processing Blades within the same BladeServer chassis for increased processing scalability.

3.2 Radar Digital Beamformer Application

Figure 4 illustrates an octal-channel FPGA-based Radar Receiver Digital Beamformer implemented on the Annapolis WILDSTAR 5 FPGA Blade. Radar beamforming techniques are used for target detection in the presence of ground clutter and jamming signals. Beamforming principles apply to both transmission and reception of signals, but this application focuses on the receive system. Figure 4 highlights the key functional blocks required to implement such a real-time processing system. Two front panel A/D mezzanine cards are used to digitize the incoming IF channels up to a sampling rate of 200 MSps. Following the A/D converter is a Digital Down Converter (DDC) for each channel, which mainly consists of parallel FIR filters. These filters split the signal into its I and Q components, reduces the sample rate of the signal, and shapes the baseband signal.

The beamforming function (Complex Multiply-Accumulate) can then be scaled across multiple Virtex-5 SXT CPE modules. In-phase and Quad-phase (real and imaginary components of a complex signal) data samples are multiplied by beam weights and the results are then summed for each beam. The WILDSTAR 5 FPGA Blade platform allows for the FPGA resources necessary to generate several beams.

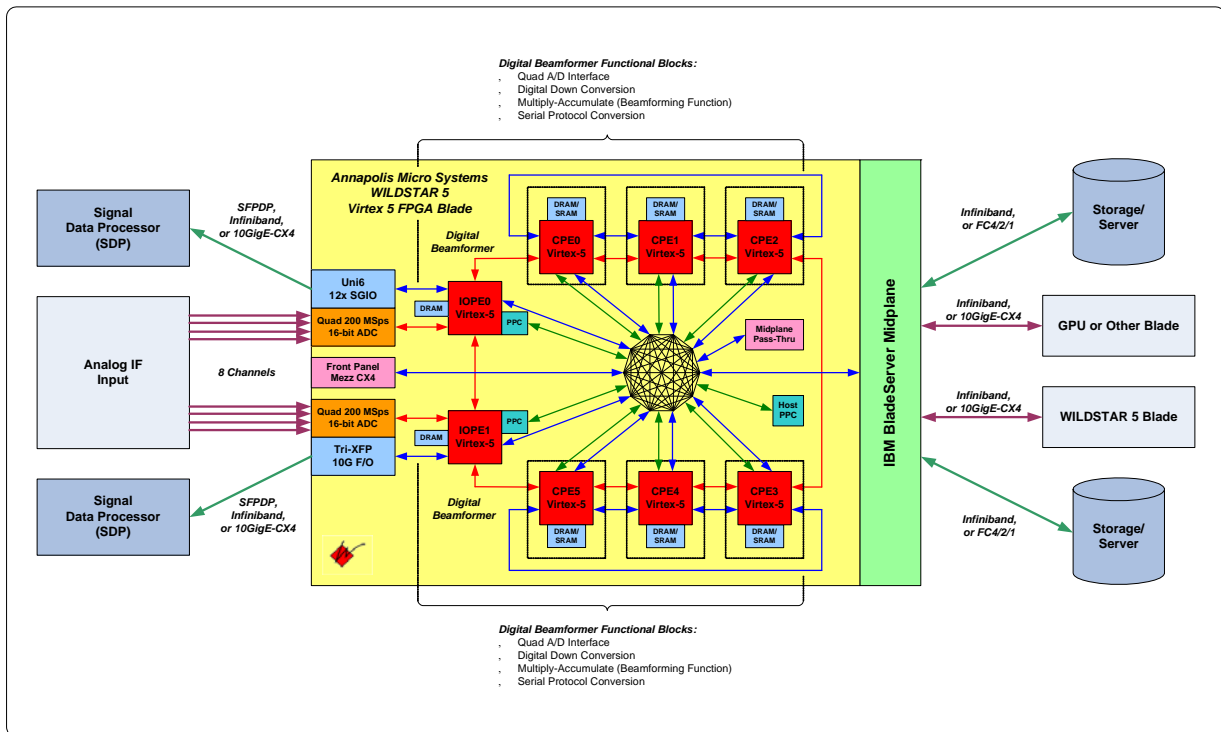


Figure 4: FPGA Blade Radar Receiver Digital Beamformer Application

Network Protocol IP cores (Serial FPDP, Infiniband or 10 Gigabit Ethernet) can then be instantiated in the IOPE FPGA to forward I and Q beam data to Signal Data Processors (SDPs) over copper or fiber media for further processing and analysis. Multiple Radar Receiver Digital Beamformer Blades can be implemented in a single IBM BladeCenter chassis to support several incoming channels from an antenna array.

3.3 Software Defined Radio Application

Figure 5 illustrates a signal processing platform for a Software Defined Radio application implemented on the Annapolis WILDSTAR 5 FPGA Blade. It highlights the key functional blocks required to implement such a real-time processing system. In the receive direction, a high-speed A/D mezzanine card is used to digitize the incoming RF channel up to a sampling rate of 2.2 GSps. Downstream to the A/D converter, a Digital Down Converter (DDC) allows for a simplification of the RF front-end design since the down-conversion process is performed in the digital domain. Digital FIR filters following the digital mixers provide much sharper filtering than those possible in the analog domain. Decimation is usually involved, thereby reducing the output data rate.

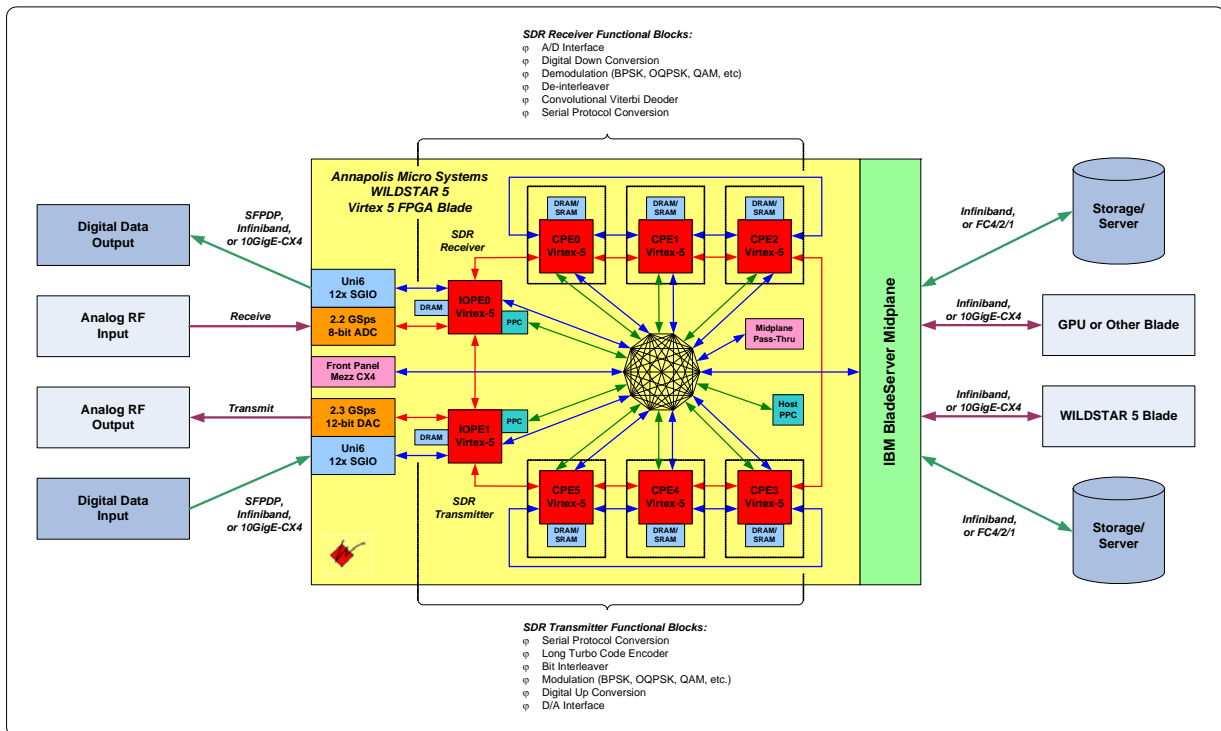


Figure 5: FPGA Blade Software Defined Radio Application

Next, virtually any form of demodulation scheme (BSK, OQPSK, QAM, OFDM, etc.) can be implemented in a Virtex-5 CPE FPGA with the appropriate algorithm. The demodulator recovers the originating signal from the IF output. Analysis functions can also be performed in this block such as energy detection which may be implemented with an FFT IP block. When new or proprietary demodulations or analysis schemes are required, no new hardware is necessary; a new DSP software algorithm is loaded into the FPGA from the host processor. Downstream receiver functions may include a Channel Equalizer, De-interleaving of frame data and Convolutional Viterbi Decoder used for Forward Error Correction (FEC), before sending packetized serial data using a standard network protocol.

In the transmit direction, the inverse functions are implemented. These include Serial Protocol Conversion, Turbo Code Encoder, Bit Interleaver, Modulator, Digital Upconverter (DUC) and finally the high-speed D/A Converter implemented on a front-panel mezzanine card. Multiple modulations schemes can be tried and proven using reconfigurable FPGA technology. The DUC uses digital interpolating filters and provides the same advantages as the DDC.

4. Scalable FPGA-based Network Processing Platform

4.1 Network Processor Application

Figure 6 illustrates a Network Processor Application implemented on the Annapolis WILDSTAR 5 IBM Blade. Ingress Network Processor (NP) functions include Serial Protocol Conversion, L4/L7 Deep Packet Classification, Policy Engine, and Packet (Payload) Buffering. Front panel Serial Gigabit I/O cards accept incoming packets, over copper or fiber media, via standard network protocols from either Multi-Protocol WAN Switches or 10 Gigabit Ethernet Switches.

Xilinx Virtex-5 FPGAs can be user programmed to perform L4 through L7 deep packet classification. Packet classification is the process of categorizing packets into “flows”. For example, all packets with the same source and destination IP addresses may be defined to form a flow. Generally, packet classification on multiple fields of the packet header is a difficult problem. Several different algorithms can be employed such as basic search algorithms, geometric algorithms, and heuristic algorithms. The algorithm suitable may be different for different types of packets. The FPGA allows for multiple classification algorithms to be programmed on a single chip.

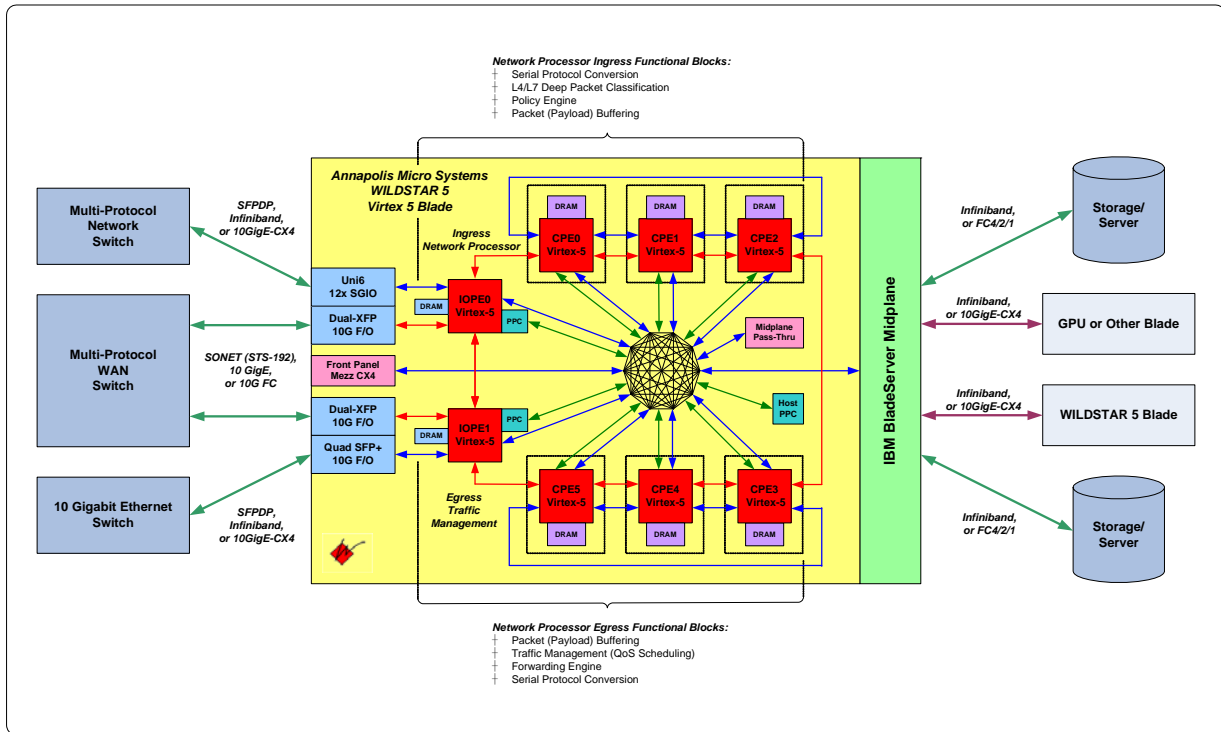


Figure 6: Blade Network Processor Application

A downstream Virtex-5 CPM can be programmed to perform Policy Control. As with Classification, there are several different types of policy control algorithms. The Virtex-5 FPGA is a good fit to implement these different policing algorithms. Such policing algorithms include Leaky-bucket, Token-bucket, and Weighted Random Early Detection (WRED) algorithms.

Egress Traffic Management functions include Traffic Management (QoS Scheduling), Forwarding Engine, and Serial Protocol Conversion. The Traffic Shaper controls the scheduling of packets and, as a result, the Quality of Service (QoS). Traffic Shaper algorithms include Weighted Fair Queuing (WFQ), Strict Priority, Round Robin (RR), Weighted Round Robin (WRR), and mixed de-queuing algorithms. Multiple Traffic Shaping algorithms can easily be implemented in a FPGA CPM. Downstream to the Traffic Shaper, the Forwarding Engine supports encapsulation, multicasting (fabric, logical, or spatial), variable size packet segmentation, packet header insertion, and flow control. Finally, Serial Protocol Conversion can be implemented in the egress IOPE FPGA, if required, before being delivering packets to external network switches.

4.2 Network Pattern Matching Application

Figure 7 illustrates a Network Pattern Matching Application implemented on the Annapolis WILDSTAR 5 IBM Blade using Virtex-5 CPMs. Pattern Matching Algorithms have many uses in Network Intrusion Detection Systems (NIDS) and IP address, security, and cyberspace surveillance systems. The difficulty in Pattern Matching is performing advanced pattern searches at line rate, i.e., 10 Gbps for each 10 Gigabit Ethernet link. Intrusion Detection analysis requires considerably more compute cycles and memory accesses per packet than required by traditional NP applications, such as IP routing and QoS scheduling. Some tasks are compute-bound and others are memory-bound. In addition, the amount of processing required for each packet is not constant.

High-speed packet content inspection devices rely on fast multi-pattern matching algorithms which are used to detect keywords or signatures in the packets. Multi-pattern matching requires intensive memory accesses which often becomes the performance bottleneck. As a result, hardware-accelerated algorithms are required for line speed packet processing. Several pattern matching techniques for NIDS have been developed for FPGAs taking advantage of its highly parallel logic resources.

Finally, the Results Processor or Decision Engine, implemented in the IOPE, uses the results of the pattern processing algorithm to take appropriate action with regards to each incoming packet.

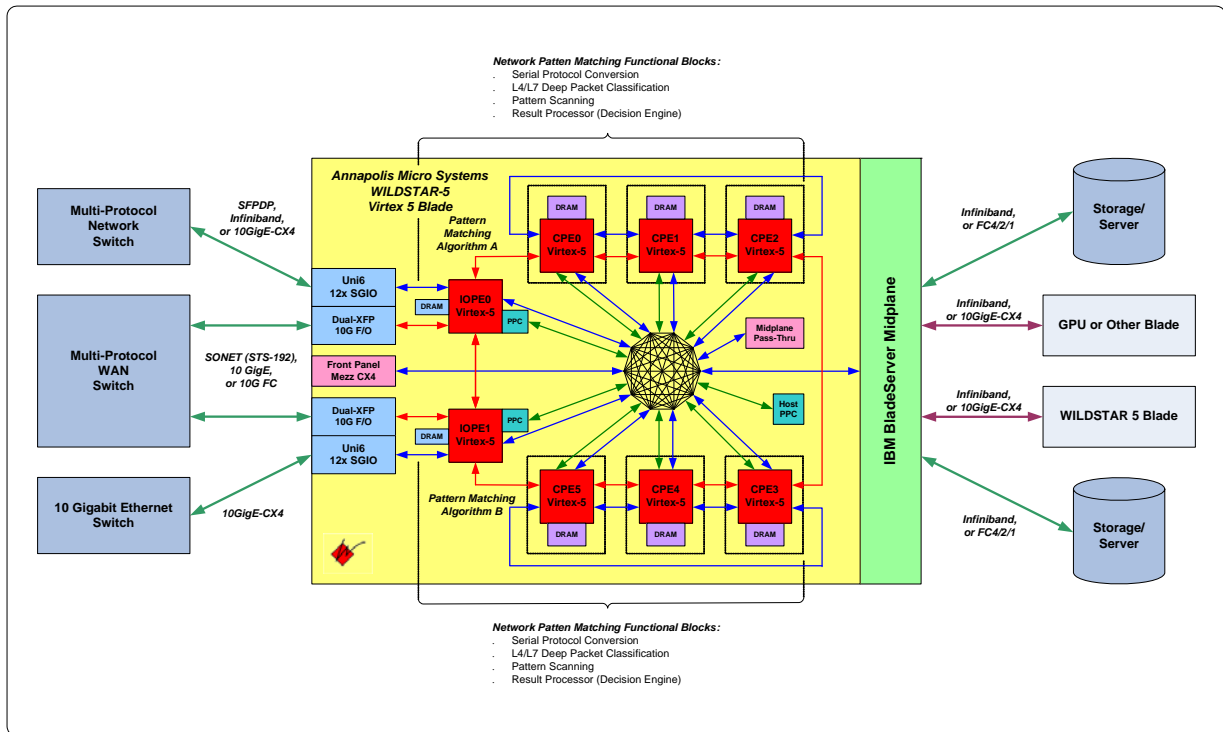


Figure 7: WILDSTAR 5 Blade Network Pattern Matching Application

6. Application Development

6.1 FPGA Development

As a high-performance FPGA computing systems provider, Annapolis Micro Systems has provided application development support using Annapolis' CoreFire™ Design Suite and traditional VHDL development environments.

6.1.1 CoreFire™ Design Suite

The CoreFire Design Suite is a dataflow-based development system used to create extremely high performance FPGA designs for the WILDSTAR 5 FPGA Blade in a fraction of the time required for a conventional VHDL-based control flow approach. CoreFire's Board Support Libraries simplify access to all the chip and board features on Annapolis FPGA products. CoreFire offers rich support for a broad range of datatypes including bit, signed and unsigned integer, floating point, variable precision floating point and arrays. It includes thousands of application level cores, supporting everything from simple adds, through DSP functions like FFTs and FIR filters, up to more complex cores like Channelizers and I/O protocol handlers.

The CoreFire Design Suite is a grouping of all the tools necessary to design, build, test, and field an application. The CoreFire Design Suite includes the following:

- CoreFire FPGA Application Builder
 - CoreFire Board Support Libraries
 - Thousands of Application Level Cores
- CoreFire FPGA Application Debugger
- CoreFire Online Help
- CoreFire Host API for Java
- CoreFire Host API for ANSI-C

CoreFire's drag-and-drop method of building designs allows for ease of use and provides system designers a simple way of visualizing designs, rather than extensive and confusing RTL code. CoreFire's patented interconnection techniques handle many internal chip complications, and guarantee design correctness over a wide range of issues. Designers from many different disciplines can use CoreFire to create FPGA applications. The CoreFire design approach consists of four steps:

1. Creation of a data flow-based design for each FPGA using the CoreFire graphical editor.
2. Place-and-Route (PAR) the FPGA data-flow design using the Xilinx ISE toolset.
3. Run and debug the design using the CoreFire Application Debugger directly on the WILDSTAR 5 FPGA Blade using the auto-generated Java or 'C' application template files. Directly access the hardware via the debug process. Start, stop, read registers and memory, step through clocks, and monitor throughput. Use your Matlab input data

file as input, and run your design on the actual hardware. Compare the output of your hardware design to the output of your Matlab simulation. This guarantees that you have implemented your simulated Matlab design on the WILDSTAR hardware exactly and completely.

4. Starting with the Java or 'C' application template files, develop your own customized host code using the CoreFire object-oriented software model that provides an interface between your user application and the runtime accessible CoreFire cores.

Annapolis' WILDSTAR™ 5 IBM Blade in conjunction with the CoreFire™ FPGA Design Suite allows for a powerful, well-integrated design environment that will reduce the time-to-market for today's complex FPGA and processor systems.

6.1.2 WILDSTAR 5 VHDL Development

The WILDSTAR 5 VHDL model includes source code and examples for all hardware interfaces. Support for Xilinx ChipScope access provides a means for real-time logic debug and verification. The WILDSTAR 5 VHDL development flow consists of six steps:

1. Creation of the VHDL design using the Annapolis supplied WILDSTAR 5 VHDL model.
2. Compilation and simulation of the VHDL design using Mentor Graphics ModelSim™.
3. Synthesis using Synplicity Synplify™ or Xilinx XST™.
4. Place-and-Route (PAR) the design using the Xilinx ISE toolset and the included Annapolis script files.
5. Develop 'C' host code using provided templates and examples to program IOPE/CPE images and run applicable WILDSTAR™ 5 Driver API function calls.
6. Run the compiled 'C' user application on the WILDSTAR™ 5 FPGA Blade platform.

7. Conclusions and Looking Ahead

Annapolis' IBM Blade product roadmap commitment will allow next generation FPGAs and Multicore technologies to be employed on today's WILDSTAR 5 IBM Blade baseboard. Other processing technologies might include Graphics Processing Units (GPUs), General Purpose or Application Specific Processors (ASPs). No other platform available today provides this flexibility.

As described by the application examples cited, customers today can leverage the BladeCenter's high speed midplane to integrate the WILDSTAR 5 for IBM BladeCenter with other products in the IBM BladeCenter ecosystem, creating a powerful and well orchestrated heterogeneous processing platform suited to today's most challenging computational tasks.

BladeCenter™ is a trademark of IBM Corporation. Virtex™ and XST™ are trademarks of Xilinx Corporation. WILDSTAR™ and CoreFire™ are trademarks of Annapolis Micro Systems, Inc. ModelSim™ is a trademark of Mentor Graphics. Synplify™ is a trademark of Synplicity.